# 5 Putting Meaning into Words

How is our static and disjointed language produced from dynamic and continuous nonverbal thought? I tried in the previous chapter to show that thought is beyond language, that meaning lies beyond language, and that it is therefore impossible to distill the essence of words into words. But there is another problem. The study of language reveals that our spoken or written words have only a partial and indirect relation to meaning and thought in any case. Nevertheless, we talk and we write, and the words we produce are quite frequently understood by ourselves and by other people. The present chapter examines how the feat might be accomplished.

## LANGUAGE—THE DEPTHS AND SURFACE

A general distinction may be made between two aspects of language, its surface structure and deep structure. The *surface structure* of language can be regarded as its physical properties—the sounds of speech, the written marks on the page for writing. *Deep structure*, on the other hand, is meaning. Most languages have more than one form of surface structure: not only speech and writing but also, for example, the gestures of sign language for the deaf and the tactile symbols of Braille. However, there is only one deep structure. No one would want to claim that the meaning of written language is in essence different from the meaning of speech.[1]

One way to look at the distinction is to regard surface structure as the part of language that exists in the world, however transiently, outside the minds of language users. We produce surface structure when we speak or write; we interpret surface structure when we listen or read.[2] Surface structure can be

counted or measured, numbers can be put to the relative loudness of speech, to the size of print, or to the rate at which words are produced. Meaning, on the other hand, defies measurement. Meaning exists in our minds—in the nonverbal, inaccessible theory of the world in our head—underlying the language we produce and making sense of the language we understand.

Meaning is not directly represented in the surface structure of language; that is the central paradox of language. Surface structure and deep structure are not reflections of each other; in fact, there is a gulf between them. There is no "one-to-one" correspondence between the surface structure and deep structure of language; every meaning can be represented by more than one surface structure, and every surface structure can have more than one meaning. The grammatical structure of a statement is not the same as the meaning structure underlying it.

That the same literal meaning can be represented in more than one surface structure is demonstrated by sentences that are paraphrases of each other. *Everyone seeks truth, truth is sought by all,* and *humanity pursues verity* are three different ways of saying the same thing, roughly at least. Conversely, sentences like *the turkey was too old to eat* and *the policeman held up the motorcyclist* each represent more than one meaning. (It may take a moment to discover two meanings for the examples I have just given, because the brain resists interpreting the same surface structure in more than one way; thus puns are not always immediately seen or invariably appreciated.) The fact that all the common words of our language have multiple meanings underlines the absence of a simple relationship between surface structure and meaning. If the word *bank* (or *time*, or *house*, or *table*—think of any common word) has more than one possible meaning, then any sentence containing that word must have more than one possible meaning. That some interpretations would not make sense is beside the point, which is that more than language is required to make sense of language. The additional factor is commonly regarded as "knowledge of the world"—a theory of the world—so that we can disentangle an anomaly like *the thieves drowned trying to reach the bank*. But the ability to make sense of language depends on more than just a theory of the world, it depends on an implicit understanding between speakers and listeners (or between writers and readers) about each other's expectations and intentions. The complex global structures of expectations and intentions are very different from the linear surface structures of sentences.

Oddly enough, although surface structure is the only aspect of language that is accessible to direct observation, the only part that can ever be counted or measured, it is usually an aspect of which we are completely unaware. We do not pay attention to individual words, and certainly not to the sounds or letters of which words are constructed, when we hear someone talk or when we read. Language is transparent; we look through the actual words to perceive the meaning beyond, in the way that we perceive a scene through a window without being aware of the window itself. We *can* become aware of



words when we listen or read, but only if we make a conscious effort, in which case we are likely to miss the meaning (just as we can focus on the glass in the window at the cost of our perception of the scene beyond). Usually we only become aware of particular words if we cannot make sense of what we are reading or hearing, when the window is obscured in some way. Meaning commands our attention.

But how is this done? How do we succeed in perceiving meaning through the ambiguous surface structures that are presented to our ears or eyes? If we cannot embed meaning in the words we speak or write, how do we manage to make ourselves understood?

A common view is that the gulf is bridged through formal grammar— that surface structure plus grammar equals meaning. But we can all understand ungrammatical sentences, for example, when small children talk. And it is not possible to say what the grammar of many words and sentences is until we understand them. We cannot say whether *bank*, *time*, *take*, *walk,* and many other words are nouns or verbs unless we make sense of them in a sentence. The single-word utterance "Fire" can be a noun or a verb, depending on whether it is shouted at a conflagration or an execution. In *open the empty bottle, open* is a verb and *empty* an adjective; in *empty the open bottle* it is the other way around. *He was seated by the usher* is a passive sentence in one sense *(the usher seated him)* but not in another *(he sat next to the usher)*. Certainly we do not produce sentences by first thinking of a grammatical structure, and then putting words of the appropriate parts of speech to that structure. Grammar is not a strong candidate to account for our ability to attribute meaning to the surface structure of language. It is not possible to get to the meaning of language from the surface down.

There is an alternative point of view that proposes what might at first glance seem an unlikely possibility, that language is understood by having meaning brought to it. We do not understand words by deriving meaning *from* them, but by bringing meaning *to* them. This perspective solves the problem of why we are not aware of the potential ambiguity of language when we speak or write. Since we normally only produce language when we have at least a general idea of what we want to say, it is not surprising that we tend to find in the language we produce the meaning we hope to express. Indeed, one of the hazards of writing is that we may know so well what we want to say that we find it difficult to examine what we have written with the less-informed eye of an independent reader. But how can we bring meaning to language when we are the listener or reader? Surely the whole point of attending to language is to discover meaning that the producer intended, not something we already know ourselves. Examination of this question will I hope throw more light on some of the operations of thought that underlie language in the first place.

During the next few pages I have a particular concern with comprehension—with the use the reader (or listener) makes of language—as well as with

composition. I have not forgotten that this book is supposed to be about writing, but one way to an understanding of what the writer does is by trying to grasp what the reader must do. After all, writers generally aim to produce something that readers can comprehend. Besides, comprehension—as I shall argue later—is the way language and its uses are learned in the first place. Composition and comprehension are in many respects inseparable.

So having failed to relate language and thought by descending from the surface to the depths, I shall now take an opposite course and consider the nature of deep structure, to see if it is any easier to move from there to the surface.

## THREE GRAMMARS OF LANGUAGE

It is not difficult to talk about the surface structure of language. We can look at the printed marks on a page of text or at transcriptions of speech and make all kinds of enumerations and analyses of words and their interrelations. All marks of grammar as it is conventionally understood reside in the surface structure of language, although there has to be some understanding of meaning before we can specify the grammatical function of particular words. We need to understand sentences to explain the purposes of their punctuation—for example, whether an apostrophe indicates a possessive or an omission. Deep structure also has a grammar, though it is different from the grammar of surface structure, and there is a grammar which links the grammars of the deep and surface structures. The surface grammar is usually the only one taught in schools. But it is the other two grammars that will be my concern for the next few pages.

### Syntactic and Semantic Grammar

What does the deep structure of language look like? The question is meaningless if taken literally, because meaning is itself an abstraction; there is no way it can be laid out for inspection as if it were surface structure; it can only be represented symbolically. Deep structure can be conceptualized in terms of elements with different kinds of relationships with each other (which is the reason it is appropriate to use the term "grammar" for the organization of deep structure), just as the elements of surface structure—the sounds, letters, and words—have varied but distinctive relationships with each other. The grammar of surface structure is *syntactic*, its relationships are concerned with the ordering or arrangement of elements. But the grammar of deep structure is *semantic*; its relationships are concerned with meaning. For example, in the deep structure of the sentence *John is playing the guitar* there are three meaningful elements, namely whatever "John," "playing," and "guitar" mean to us, and a set of meaningful relationships among them.



The deep structure of language can be represented in a number of ways, centering usually on the notion that the meaning of any sentence can be expressed in the form of a network of *propositions*, or underlying assertions. The focus of such propositions is generally an action, an event, or a state of affairs, usually but not invariably corresponding to a word that is a verb in the surface structure. For example, the deep structure of the sentence *John is playing his guitar in the garden* consists of three propositions: that John is playing, that his guitar is being played, and that the garden is the place where the playing is occurring. All of these propositions can be represented diagrammatically as follows:

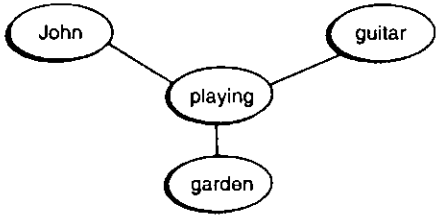


Of course, these relationships—or *"cases"* as they are known technically—are not the same. John does not have the same relationship to the playing as the guitar or the garden have. John might be called the subject of the playing, the guitar the instrument (no pun intended), and the garden the location. Linguists have identified nearly a score of common underlying relationships for English (and for all other languages, since we are now talking about the nature of *thought*, which is universal. Only surface structures differ from one language to another.)[3] Some of these relationships can be illustrated if I elaborate a little more on my exemplary sentence and deep structure: *John was enthusiastically playing Chopin mazurkas on his guitar in the garden in the rain on Tuesday to entertain his friends.*

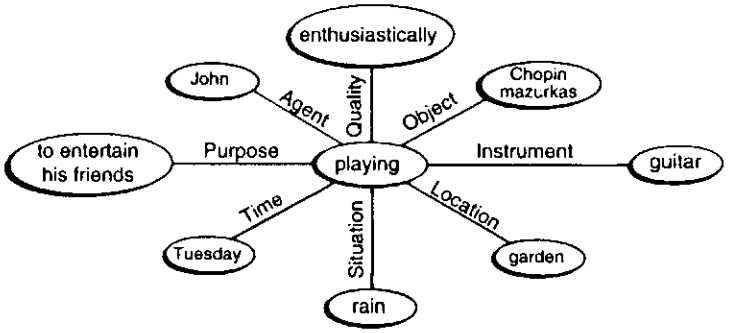

Each of the relationships can be seen as the answer to a possible question about the playing: Who was doing it? (John) How was it done? (enthusiastically) Where was it done? (in the garden) On what? When? For what reason? and so forth.

The fact that the number of deep structure elements in the examples I have given might seem to be the same as the number of major elements ("content words") in the surface structure is purely coincidental. Deep structure can and usually does contain many more elements and relationships than surface structure makes explicit, the missing parts being what the writer and reader take for granted. For example, the surface structure *John was playing his guitar on Tuesday* would usually be regarded as a reasonably complete sentence, though it leaves out, it takes for granted, any reference to the music that John was playing, where he was playing, and how and why he was playing. The sentence *A guitar was played in the garden on Tuesday* does not even specify that someone was playing the guitar.

So now we have a second grammar, a meaningful or semantic rather than a syntactic one, representing the various ways in which elements of meaning can be organized and interrelated in the deep structure of language. The two grammars are obviously not the same, or even mirror images of each other. The grammar of surface structure is *linear*, it proceeds (in written English) from left to right and the order of elements (words) makes a difference, as does their form. The grammar of deep structure does not have a sequence; thoughts are holistic, global, and no one part is prior in time to another. The semantic network remains the same, no matter how its parts might be moved around on paper. How then does the writer move from the conceptual grammar of deep structure to the syntactic grammar of surface structure? How are thoughts transformed into words? The answer is by the third grammar, a *transformational* grammar, which gives form to the elements of deep structure that the writer wishes to represent and organizes them into linear sequences of words. Transformational grammar is in theory the only grammar a writer (or speaker) actually *uses*.

## Transforming Thought into Words

Though it is not a grammar that is taught, transformational grammar is something that all speakers and writers must have, whether or not they produce language in conventionally grammatical surface structures. No one, not even an infant putting baby-talk words together in the first two-word sentences, puts words together randomly.[4] Even if some of our sentences ain't never grammatical by conventional standards, they tend to be consistently "wrong," they ain't never random arrangements of words. (I would not say "never ain't" as well as "ain't never" unless that were also a consistent



practice, a rule of my personal transformational grammar.) Transformational grammar is the link between the two other grammars, between surface structure and meaning. The other two grammars are merely descriptive; they describe the actual arrangement of word elements in a surface structure or the presumed arrangement of thought elements in a deep structure. But the transformational grammar is productive; it *generates* surface structures from meaning. Drawing upon an inner dictionary or "lexicon" of its own, our personal transformational grammar selects lexical items which are appropriate to the deep structure elements which are to be expressed, and modifies and organizes these elements into what we can recognize as a surface structure of words in speech or writing.

Transformational grammars employ three principal devices to represent in surface structure some of the meaningful elements and interrelationships of deep structure. The first device is word order. In English, usually, the agent or subject of a sentence tends to be at the beginning. Thus, *John is playing the guitar* does not mean the same as *the guitar is playing John*. The order of words in a sentence is a crucial part of English, although it does not make a difference in other languages which rely more on the second transformational device for representing deep structure interrelationships, namely variation in the form of a word. The change in the form of a noun depending on its meaningful function in a sentence is the sense in which the word *case* is usually used in formal grammar, for example in referring to heavily inflected languages like Latin, Greek, or German. In such languages the form of a word like *John* or *guitar* would change depending on whether it was the subject or object of a sentence, while in English only pronouns change depending on the subject-object distinction (we say *he saw her* but *she was seen by him*). The third device for representing underlying case relationships is employed much more in English than in many other languages, namely the use of special words to express relationships. The special words in English (and many other European languages) are primarily prepositions such as *in*, *at*, *by*, *with*, *to*, and so forth, small ubiquitous words with a wide variety of uses, making them the hardest part of language to learn or to translate.

For some reason, which has never been adequately explained or even (I think) considered, there is a remarkable complexity among the prepositions that represent various case relationships. The same prepositions can be used to represent quite different kinds of relationship; I can say I hope to be *on* form *on* Friday when I shall be *on* the platform *on* time to talk *on* linguistics to a group *on* vacation from their courses *on* education—seven different uses of the preposition *on*. We also use different prepositions to express similar kinds of relationship. For example, we may say there was a party *at* someone's home, or *in* their house, or *on* their premises (*under* their auspices, *with* their approval, *by* their leave).

Because the same preposition can be used for quite different case relationships, the preposition alone cannot be taken as an indication of the underlying case relationship; this is a major cause of the ambiguity of the surface structure of English. For example, John may have been playing Chopin on his guitar on Tuesday, but the two instances of the word *on* represent different case relationships; they have different meanings, as we see if we try to join the phrases they are in with *and*. We cannot say John was playing Chopin on his guitar and Wednesday, any more than we can say he was playing in the rain and his best suit, or that we plan to send him an invitation by mail and Friday. But prepositional phrases with the same underlying meaning can be connected by *and*, so that we can say that John was in his blue shirt and white shoes, or entertaining on his flute and guitar. Sometimes the same surface structure word order will similarly represent quite different underlying meanings, as when we say that John was cooking at the barbecue and that hamburgers were cooking at the barbecue, but again the "and" test for similarity of underlying relationships applies and we find we cannot say that John and hamburgers were cooking at the barbecue (unless they were indeed both cooking in the same sense). The similarity of surface structures for different underlying meanings is a linguistic coincidence that writers can generally ignore.

## Transforming Words into Thought

But what about comprehension? Will transformational grammar work in reverse for readers (or listeners), and why are they not usually aware of the potential ambiguity of the surface structures of language? The answer is that transformational grammar does not work in the reverse direction, and that readers must understand writing by employing their own transformational grammars in the same direction as the writer, proceeding from the deep structure to the surface.

At first sight it might seem absurd to suggest that readers are themselves responsible for bringing meaning to what we write, for imposing sense upon the printed marks on the page. But all that this assertion entails is that readers understand what the author is talking about, which does not I hope sound such an enormous or mystifying requirement. It is obvious that a reader who does not already know something about the language the author is using and the subject matter the author is dealing with will not make sense of what the author writes. By bringing to bear prior knowledge of the author's language and subject matter, the reader can anticipate possible surface structures on the page and thus find sense in them. The writer must anticipate what the reader will expect. I am not arguing that a reader should be able to forecast exactly what the author will say. If the reader has that much prior knowledge, there would be no point in reading in the first place. But the reader must always anticipate what the writer *might* say, in order to understand what the author *does* say.



It is not difficult to demonstrate that prediction is a constant part of the way we understand language (although it is again not a part of which we are usually aware); indeed, prediction is a constant part of the way we understand the world in general. The demonstration is that it is always possible to be surprised. I could write something that surprises you—by suddenly changing the subject matter to a discussion of scuba diving, or by switching the language from English to French, or by doing a variety of other idiosyncratic things. It is always possible that something will surprise you as you drive your car downtown—it is conceivable that you might see a giraffe or a deep sea diver in the main square, or you might meet your Uncle Fred who you thought had taken up permanent residence in Nepal. But events like these would surprise you for the very reason that you did not anticipate them. That is what a surprise is, something that is unexpected. The fact that we are so rarely surprised is not because we usually expect that anything might happen, but because our expectations are usually very accurate predictions of what indeed does happen (or rather, of what might happen).

I suspect that it is the reader's need to predict that explains and even justifies the importance frequently attributed to formal grammar, the rules of usage that are taught at schools and in the grammar books. Grammar is not essential for comprehension. If we can detect errors in the grammar of a sentence and put them right, then we do not need the grammar to tell us what the sentence intends in the first place. But writing that is ungrammatical confounds our expectations. Like unconventional spelling and punctuation, it makes writing more difficult to anticipate, although it is rarely the cause of its being incomprehensible.

## What Transformational Grammar Lacks

Transformational grammar is a theory that endeavors to describe hypothetically the manner in which language and thought are related. The theory is entirely conjectural, of course; there is no means by which we can look inside the brain to see if such a grammar is really there. But the grammar is a useful metaphor, a way of thinking about what might otherwise be completely unthinkable and indescribable, the actual processes of the brain. It has its limitations, nonetheless.

The theory is simplistic, for a start. Transformational grammar as I have outlined it can only produce simple declarative sentences—nothing like the complicated compound constructions that the present sentence exemplifies, nor such alternative sentence structures as interrogatives or imperatives. But these objections can be dealt with handily. Additional transformational rules can be proposed for combining two or more simple sentences into one

compound one, or for generating alternative sentence structures. More serious are objections that transformational grammar, no matter how well it might appear to relate the surface and deep structures of language, still is a very unrealistic model of the way in which language production and comprehension actually take place. As a working model, transformational grammar has a certain lack.

Consider writing. While logically transformational grammar might seem able to generate surface structures from meaning, the entire mechanism looks remarkably inert. What puts it to work? What determines that one transformational rule rather than another will be employed? Where exactly do deep structures come from, what do semantic networks represent, how does a writer move from one deep structure to another, and what decides exactly which deep structure elements will be represented in surface structure? Where is the dynamic element to put some force and direction into it all?

Consider reading. It is tempting to believe that an appropriate semantic network (like the guitar-playing example on page 53) is the meaning of the sentence, that a sentence has been understood if it has been reduced to deep structure elements and relationships. Many language researchers seem to assume that this is the case, and they write computer programs that are considered to simulate comprehension if they can reproduce underlying semantic structures or can answer questions about the surface structure. But comprehension is not paraphrase. We do not usually understand what we read or hear by translating it into other words, or even into some abstract or symbolic form. Something even deeper would then be required to interpret the deep structure. Where again is the direction, the reason behind the transformation?

Something is lacking in transformational grammar from both the writer's and the reader's point of view. And that lack has something to do with the dynamics underlying language, with purpose and direction. What is lacking is *intention*.

Because we intend to say something, one deep structure rather than another is selected. Deep structures are part of our theory of the world generated and shaped by intentions. Because we want to write or say a particular thing to a particular audience in particular circumstances, one way of saying it rather than another is selected. We do not blindly express everything that is in a deep structure, but only that which we believe our reader or listener needs to know in order to understand or respond appropriately to our intention. Once one thing is said, we become ready to say another, or to stop. We do not repeat unless we have a reason. Intentions provide continuity as well as purpose. Each new step along the way is determined with respect to the point we have just reached, where we might get to next and our more general, global intentions.



Intentions underlie comprehension as well. We never attend to language for no reason at all. When we read, we have a purpose. We read with certain expectations. Our purpose reflects global intentions and our expectations generate focal intentions (to look for particular things); there is nothing passive about reading or listening. We do not read a sentence in order to generate a deep structure; we read it *from* a deep structure, seeking to fill gaps, to answer questions, to confirm expectations, and to reduce uncertainty among alternatives. The deep structures of comprehension are no more static than those of writing. We move from one to another—or rather our deep structures merge one into another—as our thought moves through the text we are reading. It is true that our intentions as we read might seem to develop along lines determined by the text, by the author. But unless we have intentions to explore the text in a particular way (constrained by the text in just the same way that our intention to explore a river is constrained by the actual course of that river), then there can be no way for the author to capture our attention in the first place. The complex and intimate relationship between a writer's intentions and a reader's expectations—the manner in which each influences the other across time and space and through the barrier of a solid wall of text—will be the topic of Chapter 7.

## THE CONVENTIONS OF LANGUAGE

The metaphorical manner in which I have talked of transformational grammar—as if it were a machine, a "device for producing sentences" or a set of procedures by which grammatical and meaningful sentences are generated—is not uncommon. These are all useful analogies, but they can be misleading. They are especially misleading if they are taken to suggest that there is some kind of logical necessity about the particular ways in which the surface structures of language are constructed, reinforcing perhaps an already widespread assumption that certain forms of language have an inherent "rightness" about them.

I have an alternative point of view—that every aspect of language is permeated by conventions which are in essence arbitrary and accidental. There is no inevitable logic about the particular forms of grammar, of spelling, of punctuation, or any other aspect of language apart from the fact that every convention serves a purpose; it is useful. Every convention of language could be changed and still function as effectively, provided the changed form became accepted as the convention. There is no inherent "rightness" about any aspect of language, only accepted or imposed custom. This is the reason I prefer to talk about grammar, spelling, and punctuation in terms of "conventional" and "unconventional" rather than "correct" or "incorrect." In lan-

guage, as in every other form of human behavior, it is the conventional that tends to be regarded as correct, logical, and inevitable, especially by those who themselves observe the convention. The surface structures produced by our transformational grammars become grammatical and meaningful to the extent that they conform to convention, not because they contain some intrinsic appropriateness or relationship to meaning. Conventions work because they exist, and totally different conventions would work just as well—"yes" could mean "no" and "no" mean "yes"—provided we all agreed to use words in the new way.

## The Utility of Conventions

By "convention" I mean simply an expected way of doing something that is accepted by all parties concerned, an implicit understanding that there is just one most appropriate way of doing something, not because that way is necessarily best or obvious or even particularly logical but because absence of a convention or disagreement about what it should be would cause inconvenience if not confusion. Conventions are so important and pervasive in human affairs that they often have the force of law behind them. But such is their importance and pervasiveness that generally they arise and persist without any need for legal enforcement at all.[5]

Take, as a general example, the fact that in every culture that has roads and traffic it is conventional for vehicles to drive on one side of the road rather than the other. It does not matter whether all vehicles drive on the right or on the left—different regions of the world have different conventions—provided that everyone in a particular region drives and expects everyone else to drive on the same conventional side. There is no logic behind the particular convention except consistency, established practice, or historical precedent—in other words, conformity with a convention that already exists. No committee ever sat down to debate whether there was any *essential* advantage to driving on the left (or right) rather than the other. The only possible consideration is conformity with other conventions. (The fact that vehicles might already have their steering mechanisms on one particular side is an existing convention.) You can predict that traffic everywhere in the world will be driven consistently on one side of the road rather than the other, but unless you have relevant historical information you will be unable to predict what that side will be. Reasoning will not help because there is no particular reason, only the historical accident of custom.

Here is another example, closer to language. Every literate culture that has an alphabet also has a conventional sequence in which the letters of that alphabet are arranged, an alphabetical order. I can find no logical reason for the particular order that exists for the English or any other alphabet except



tradition and historical precedent. Alphabetical order is a convention, and is respected as such. There is certainly a logic in having an alphabetical order, without which dictionaries, directories, libraries, and all bureaucracies would be thrown into complete disarray. But there is no particular logic in the order itself. Any other order would work just as well provided the new order became the conventional one. Some logic could be introduced into the order; letters could be arranged in order of relative frequency as printers sometimes organize them, or all the vowels could come first, or thin letters could be at the beginning and fat letters at the end. Any logic would make the order more predictable. But there is nothing but historical knowledge of earlier conventions that would make our alphabetical order or any other predictable. Just as there is no intrinsic logic in alphabetical order in the first place, so there would be no particular logic in changing it. Alphabetical order is useful because it exists, and to try to change the convention would undermine its utility. Spelling reformers have always had limited success because they try to pit logic against convention and underestimate the utility of convention itself. It is convention, not habit, that is hard to change.

I am not saying that language itself is illogical or irrational. There is a good reason behind every convention in language. But that good reason never explains the particular convention, only the existence of convention. I am also not saying that there is not consistency in language. With various self-evident exceptions, English forms its plurals in consistent ways, just as there are consistent ways of forming the past tense. But plurals and tenses are formed in other ways in other languages, so there is no particular logical necessity about the English conventions, and knowledge of how the plural is formed in English (or any other language) will not help one predict how the different tenses will be formed. Consistency within a convention does not permit prediction beyond the convention; only a prior knowledge of the kinds of conventions to which a language adheres allows one to anticipate what a convention might be or what it might mean.

Why emphasize that the forms of language are conventional, established arbitrarily or by chance rather than on the basis of reason or intrinsic logic? For the important reason that language can never be produced or understood on the basis of its own intrinsic reason or logic; language does not explain itself. Writers and readers must share the conventions they employ and encounter; they cannot work out what a convention is likely to be. To express a conventional meaning a writer must embed it in a conventional form; there is no "logical" or "rational" form for any meaning to take beyond consistency with the conventions of the language itself. Readers must know the conventions too; they cannot work out (without clues from an understood context) the meaning that an unfamiliar convention is intended to express; they cannot *deduce* a meaning or intention from surface structure itself.

The arbitrary nature of language is also critical when considering how ability in writing or any other aspect of language is developed. Aspiring writers cannot predict or work out what particular forms of written language are likely to be; rather, they must discover what the conventions are. The conventions that are language must be experienced and *learned.*

## The Pervasiveness of Conventions

Language consists of a truly staggering number of conventions. There is a convention for everything that language does. Writing has all its conventions for spelling, for punctuation, for paragraphing and capitalization, just as spoken language has its unique conventions for the sounds of words and for intonation. No two languages even share the same set of sounds from which their words can be constructed.[6] There are conventions for the words themselves, for their "names," and for their meanings. There is no logical reason (apart from conformity to an earlier convention) why a particular object should be called "table" in English or "mesa" in Spanish, or why these words have the particular range of different meanings that they do in each language. There are conventions for the ways in which words may be arranged grammatically and meaningfully into sentences and conventions about how sentences themselves are interrelated. In spoken language there are complicated yet extensively employed and understood conventions for taking turns in conversation and for other language interactions, for who can interrupt, and when and how to interrupt, and how to ignore an interruption. And all of these conventions vary in conventional ways depending on who is talking (or writing) to whom, their relative status, and other circumstances. How I ask someone to pass the salt depends on whether I am making the request at home or in a restaurant; whether I am asking a friend, a stranger, or a server; on our relative ages and the various relationships that we might have to each other; whether one or other of us has an arm in a sling or whether there is some other reason for my not being willing or able to reach for the salt myself. There are conventions for everything—easily demonstrated because there is always a multitude of wrong ways of saying anything, in speech or in writing.

Just as spoken language conventionally differs depending on the situation in which people are speaking, so different conventions of writing must be respected according to the circumstances that writers and readers are in, on the role of the text. These clusters of conventions are known as *genres.* Personal correspondence is one written language genre and business correspondence another—you do not write in the same way for both. Newspapers are a different genre (or rather, set of genres) from magazines, and articles must be written differently depending on the journal they are intended for. Stories constitute different genres from company reports, and both are



different from poetry and plays. There are countless written language genres, which have arisen—like the conventions of which they are comprised—in countless arbitrary and accidental ways. Newspapers in Europe do not *look* like newspapers in North America, not because of any fundamental differences between European and North American writers or readers, but because of a myriad of unpredictable events in the evolution of newspapers on both sides of the Atlantic. There is no way of *figuring out* what the conventions of any particular genre might be, and very little chance of having it *explained*; the only way to gain a working acquaintance with the conventions of a language genre is to participate in that genre, which for the genres of writing means *reading.*

The term genre originally referred to broad different categories of writing, such as comedy, tragedy, epic. More recently the term has been used to refer to the structure and organization of different kinds of media—newspapers, periodicals, novels, "nonfiction"—all of which have their own conventions. Current emphasis is on complete settings—the language and the total situation in which it is embedded—and includes conversations, interrogations, and classroom procedures. The teaching of genre has been the center of a notable educational controversy in Australia, as I describe in Chapter 14.

There are even conventions for every nonverbal aspect of language.[7] There are conventions for how long you may look people in the eyes when you talk to them, how close you may stand to them, how much you may touch them, and where you may touch them, and again these conventions—the "appropriate behavior"—vary depending upon a complex range of factors including the number of people involved, their relative ages and status, the formality of the occasion, and their intentions and expectations.

It is entirely predictable that there will be conventions for all these aspects of language in every culture in the world. But it is entirely unpredictable what the particular conventions will be; there is no way for a visitor to an unfamiliar country or culture to anticipate the form of any convention. The mark of respect may be to gaze steadfastly at the other person or to avert one's eyes, to be higher than the other person or lower, to stand close to the other person or to keep distant. Every movement has a meaning, an implication; there is a convention for everything and every convention serves a purpose. Yet what makes the convention work is simply that it exists, not what it happens to be. Every aspect of language that we take for granted is a convention that we must have learned. Every language user's knowledge of language is enormous. Yet we acquire this knowledge—for the language to which we are accustomed—so easily and imperceptibly that we are rarely aware that the conventions exist, let alone that we have learned them.

There is no completely free choice in language; there is a convention for everything, and every convention makes a difference. I mean that there are never two entirely substitutable ways of saying the same thing. Words cannot

be selected and organized idiosyncratically. Say something in one way and it has one set of implications; say it with even the slightest of changes and something different is said. The differences are often most obvious when the change is slightest, saying "Would you pass the salt" rather than "Could you pass the salt," for example. If alternatives exist, then they serve different purposes. While different surface structures may represent the same literal meaning (as I showed earlier in this chapter), they are never substitutable in practice. One form is always more appropriate or preferred than another in a particular context. The choice does not depend upon the intrinsic meaning of the sentence, on the proposition that it expresses, but on the purpose for which it is produced.[8]

## Conventions and Creativity

The idea that language affords no freedom offends many people (as it unsettled me when I first considered it). I am not saying that we are not free to say anything we want to say—a different question altogether—but that given that we have something to say, there will be only one conventional way of saying it on any particular occasion. Where we feel we have a choice, it is always between saying two different things, however slight the difference, rather than between alternative ways of saying the same thing. There are no identical alternatives in language, no paraphrases, no synonyms. *The dog chases the ball* may represent the same literal meaning as *the ball is chased by the dog*, but in a meaningful context or setting the two phrases are not substitutable. The first phrase is a statement about the dog, the second a statement about the ball. The terms *bachelor* and *unmarried man* may appear to have the same meaning, but one term is always more appropriate in a given situation. If I tell you my brother is an unmarried man you will wonder if I think you do not know the meaning of "bachelor." If I say "petrol" rather than "gas," "biscuit" rather than "cookie," I reveal something about my own language background or about my perception of the language environment that I am in.

It may also be objected that if everything in language is conventional, if there is a convention for everything and no freedom of choice, then how can anything original ever be said? How can we find language to express something new? But paradoxically, it is because conventions exist that new things can be said. Without convention, how could one ever be unconventional? Creative writers deliberately contravene convention to make a point. James Joyce flouted conventional grammar and e. e. cummings disregarded conventional punctuation, just as Picasso rejected conventions of naturalistic painting. Artists and innovators of every kind use convention by selectively ignoring it. Their unconventionality expresses new meanings, different purposes. But they would be unable to make their statements if there were no conventions in the first place. Of course, the innovator runs risks. If unconventionality is attributed to ignorance or intransigence or if its purpose remains opaque, the creativity will fail. Unconventionality only works when its purpose is understood. And when a contravention of convention is particularly successful, when its utility becomes evident to a number of people, then it may well become conventional itself. Thus new conventions arise and old conventions die.



Conventions can always be broken; we are always free to express something unconventionally, though not if we wish to convey a conventional meaning. Conventions may always be broken for good reason, and there are basically two good reasons for conventions to be broken. The first reason is ignorance, and the second is to say something new. It is always permissible to break a convention through ignorance—provided we are prepared to be characterized as a person not familiar with the convention. Foreigners are expected to break convention—that is a condition of their foreignness. If they demonstrate knowledge of relevant conventions then they are not regarded as foreign—and unconventional behavior is regarded differently. Sometimes it is convenient to be "foreign," because different behavior may then be excused. But when convention is contravened and ignorance is not an explanation, it is always assumed that the contravention is for a purpose, that the contravener has something particular to say.

Metaphors are deliberate contraventions of the conventional use of words. It has often been argued that it is only through metaphor that we can get new ideas into the brain. The first automobiles were horseless carriages. Certainly children learn by using language metaphorically. A cow is a horse with horns and the moon a lamp in the sky. Metaphors lose their force when their meaning becomes conventional—like *dead ends* and *frozen metaphors*. Many of today's most common words and expressions began as metaphors; their utility has made them conventional.[9]

Why have conventions? Why not allow language to be creative all the time? I have argued that the existence of conventions permits new things to be said, and also alluded to the important consideration that different conventions permit us to claim status and show membership of particular social and cultural groups. But there are two more fundamental reasons why conventions develop in the first place. The first is simply the nature of the human brain. Fortuitously or not, the brain does not care for alternatives that do not have a distinct purpose. Variation is expected to be motivated. Perhaps the brain has enough difficulty trying to find one meaning, one explanation, for every phenomenon it experiences in the world, without the complication of differences that do not make a difference. The brain likes to predict what will happen in every circumstance. This leads to the second basic reason for convention: It makes communication possible.

### Bridging the Language Gulf

I have tried to show that a gulf always exists between the surface structure and deep structure of language, that the basis of language comprehension is prediction. If listeners and readers are to understand language, then they must anticipate the surface structures that speakers and writers are likely to produce. And, conversely, speakers and writers must produce the kinds of surface structure that listeners and readers will anticipate. For all their complexity, the languages of the world are not fundamentally different from the kind of mutual mind reading that takes place when two entirely isolated prisoners try to communicate by tapping on the cell wall between them. Each has to imagine what the other would be most likely to want to say and how the other would be most likely to try to say it.[10]

Conventions permit static sequences of words to convey information and create images of a world in flux. Words can never directly represent either continuous changes of state that occur with the passage of time, or the simultaneity of all aspects of a simple event.[11] But if I say "The door was opened suddenly" in circumstances in which you would use the same words, then I can in fact describe the movement to you. If I put together the words "John was playing the guitar in the rain" in a conventional order then you can reconstruct an event that did not have a sequential order but was an indivisible whole. What matters is not the convention, but the implicit agreement behind the convention. Give anything a name and you can talk about it with anyone else who gives the same thing the same name. Agree on the significance of a descriptive label, and objects or events can be described in terms of that label. Every infant who learns language makes an implicit contract to use and interpret language in the terms in which the infant perceives people around using and interpreting language. It is all tapping on the cell walls but on an incredibly subtle and complex scale.

Language works to the extent that it is predictable. But for readers to predict they must anticipate the conventions that writers will observe (or contravene), and writers must respect (or carefully contravene) conventions that readers will expect. It may seem paradoxical that the conventions themselves should be quite unpredictable, that they need to be learned in advance, but it is due to its essential arbitrariness that language succeeds in conveying meaning, in making sense, and in constructing worlds in which different people may have similar experiences.

## THE INTERACTION OF THOUGHT AND LANGUAGE

To recapitulate briefly from this and the previous chapter: Thought is essentially nonverbal. It proceeds in its complex and unobservable ways in and through the theory of the world that is contained within our brains. We



cannot inspect thought directly any more than we can directly inspect what we know, but we can make our thought and knowledge manifest. We can observe the products of thought, through the conventions of language and through other kinds of behavior, either overt or imagined, just as we can observe the language of others and other kinds of occurrences in the world around us.

Our language is generated by our thought, but it is not itself thought, to which it is arbitrarily and conventionally related. Language is a product of thought, but not its image. Language is related to thought the way cooking is related to thought; both are products of the human brain and both reflect the way we think. Neither is an *image* of underlying brain processes. We should no more expect to find sentences in the brain than we should expect to find soups or casseroles there. But while language (like cooking) owes its existence to thought, it also *influences* our thought. Even when we keep our language to ourselves, either in our own minds or in writing which is not shown to anyone else, the language affects how we think. There is an interaction between language and thought.[12]

First, the language that our thought produces modifies our thought as it is produced, just as the image in a painter's mind develops as each brush stroke is applied to the canvas. The brush stroke is never in the painter's mind until it is produced or imagined; until that moment there is just a generalized intention, as intangible as the intention that sparks a word. Language permits thought to fold back on itself; the product of thought itself becomes an object that thought can operate upon, like the painter's brush mark, and thereby provides a basis for new or modified ideas. We can contemplate a statement we have ourselves made just as we contemplate the statements of others, and go forward to make further statements that would not have been possible if the original statement had not been contemplated. In the same way that a painter reflecting upon an unfinished picture can move ahead in new ways that would not have been possible—or conceivable—before the picture had progressed so far, so does language offer ways to create and explore new worlds.

Language is able to influence thought because it is able to—in fact, it has to—freeze time. Language slices a continuous progression into separate events, into instants of time. Thought is always ongoing and undifferentiated; it is not a sequence of separate states or acts. To examine itself thought must halt its progression through time in the same way that a still photograph stops action. Language can only be manifested in terms of concurrent and successive events, both in its own structure and in its descriptions. So what is a limitation when language is used to describe thought or the world—that it must arbitrarily categorize and partition continuities into events—becomes an advantage when thought tries to become reflective. Language gives thought something different to consider.

This constant possibility of interaction between thought and the surface

structure of language—each remaining apart yet powerfully capable of influencing the other—is perhaps best represented in a diagram:

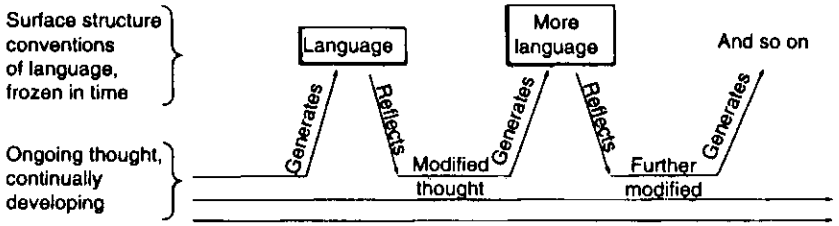


The diagram could be taken as a representation of the manner in which thought responds to language produced by others—the top part representing a speaker or writer, the lower part a listener or reader. But we can read what we ourselves have written, hear what we ourselves say, and respond to the language we produce as if it were something new because in fact our own language can be something new. What we say is created out of our minds but was never part of our minds. The conventions of language can create a new world—or at least a different world—for ourselves as well as for other people.

This potential of language to construct new worlds is one reason I regret the tendency, especially widespread in education, to regard language primarily as *communication*, as a vehicle for transmitting *information* from one person to another. How can one communicate anything new to oneself? How can you inform yourself of something you did not know already? The perspective completely ignores the possibility that language can be the means of creating worlds and of exploring ideas, our own as well as other people's, that language creates as well as communicates.

And for the exploration of our own ideas, writing has particular advantages over speech and over imagined language as well. It is to the differences between writing and speech that I now turn.

(Notes to Chapter 5 begin on page 245.)

# 6 Language: Spoken and Written

It is not difficult to see, or to hear, that written and spoken language are not the same. The written transcript of a spontaneous speech reads quite differently from a text written to be read, and we can always detect whether a speaker is talking extemporaneously or reading aloud from a text.

The exact differences between written and spoken language are many and complex. They are also subtle. It is not that written language has one vocabulary and grammar, and spoken language another, but that the vocabulary and grammar they share are used in different ways. They employ different conventions. And they differ, I argue, for good reasons.

## WHY WRITTEN AND SPOKEN LANGUAGE SHOULD BE DIFFERENT

Take first the point of view of the language producer, the writer or the speaker. Why should language that flows from the fingers be different from that emerging from the mouth?

One obvious reason for a difference is that speakers are likely to get more direct and immediate reactions from their intended audience than writers. If I talk to you, I can see whether you understand, whether you are puzzled or bored, involved or impatient, behind me or ahead. You signal to me whether I should repeat a remark, add a detail, or hurry along and skip a point or two that I was intending to make. Less obviously, the mere fact of your attention provides an incentive for me to keep talking, and I am unlikely to grind to a halt in the middle of a statement doubting whether I am being heard at all.