8

Meaning and thought

Does the language we speak influence the way we think? Does it even render certain thoughts unthinkable? Or guide us inexorably towards certain others?

This view, sometimes in a very strong form, seems to have a certain currency. This is partly through its endorsement by various of the more fashionable philosophers and literary theorists. Ludwig Wittgenstein (1889–1951), for example, opined, 'The limits of my language mean the limits of my world.' The influence of the view can also be attributed to George Orwell's *Nineteen Eighty-Four*, in which totalitarian philologists employed by the Party have concocted Newspeak, a language which is supposed to make its speakers incapable of expressing or even entertaining any thought that deviates from the orthodoxy of Ingsoc (English Socialism). Here is how Orwell describes the language in the Appendix to his novel:

> The purpose of Newspeak was not only to provide a medium of expression for the world-view and mental habits proper to the devotees of Ingsoc, but to make all other modes of thought impossible. It was intended that when Newspeak had been adopted once and for all and Oldspeak forgotten, a heretical thought—that is, a thought diverging from the principles of Ingsoc—should be literally unthinkable, at least so far as thought is dependent on words.

Although this is strong stuff, expressing a remarkable confidence in the ability of language to shape thought, it is worth noting the loophole 'at least so far as thought is dependent on words'. Orwell

does not quite close the door to thought not being wholly dependent on words.

As well as Orwell, the potentially stultifying effects of our language are particularly associated with the names of the linguists Edward Sapir (1884–1939) and Benjamin Lee Whorf (1897–1941). Sapir was born in Germany but emigrated to the USA, where he taught anthropology at Yale; Whorf was American, by trade a fire inspector for an insurance company, but a talented linguist and a student of Sapir's. The two are associated with what is now called the *Sapir-Whorf hypothesis*, though they did not coin this term themselves. But what exactly was the hypothesis? Sapir and Whorf's writings on this matter tend to be vivid but vague, as the following quotation from Whorf illustrates:

> We dissect nature along lines laid down by our native languages. The categories and types that we isolate from the world of phenomena we do not find there because they stare every observer in the face; on the contrary, the world is presented in a kaleidoscopic flux of impressions which has to be organized by our minds—and this means largely by the linguistic systems in our minds. We cut nature up, organize it into concepts, and ascribe significances as we do, largely because we are parties to an agreement to organize it in this way—an agreement that holds throughout our speech community and is codified in the patterns of our language. The agreement is, of course, an implicit and unstated one, *but its terms are absolutely obligatory*; we cannot talk at all except by subscribing to the organization and classification of data which the agreement decrees.

This is all very portentous. The basic thought seems to be that the only conceptual distinctions we can make are those encoded in our language; and that the reason for this is that our language imposes those distinctions on the sense data we take in from the world. But when we look more closely, we become unsure about what exactly is being claimed. We read that the kaleidoscopic flux of impressions with which we are buffeted is organized *largely* by the linguistic systems in our minds; that we cut nature up as we do *largely* because we are parties to a linguistically sustained agreement along those lines.

The exact reach of these largelys is not specified. Similar hedges can be found in Sapir's writing on the subject.

But we should not judge a view solely on the basis of people's unwillingness to express it. I will here present three versions of the Sapir-Whorf hypothesis, including a very strong one, based on Sapir and Whorf's own writings and on related work in the current technical literature:

(1) a. The strong Sapir-Whorf hypothesis. The only conceptual distinctions we can make are those encoded in our language; and the reason for this is that our language imposes those distinctions on our sense data (and we have no other source for such distinctions).

b. The restricted Sapir-Whorf hypothesis. There are some topics such that the only conceptual distinctions we can make regarding them are those encoded in our language; and the reason for this is that our language imposes those distinctions on the relevant sense data (and we have no other source for such distinctions).

c. The watered-down Sapir-Whorf hypothesis. There are some topics such that the way we habitually or stereotypically think about them is influenced by the language we speak.

The strong Sapir-Whorf hypothesis is based on the passage of Whorf quoted above, with the adverb *largely* removed. The restricted version keeps the idea that we are incapable of making conceptual distinctions that are not made in our language, but allows that this might hold only with respect to certain topics. And the watered-down version gives up on the central idea of the first two versions and allows that we might, after all, be able to make conceptual distinctions not provided by our language; but it maintains, nevertheless, that our language does have some less drastic influence on our thought, by influencing the way that we habitually or stereotypically think about some topics. If this last hypothesis is correct, we will naturally want to ask how many such topics there are and to what extent we are influenced in this way.

The strong Sapir-Whorf hypothesis, then, says that the only conceptual distinctions we can make are those encoded in our language. It is universally agreed among cognitive scientists that this hypothesis is false. For one thing, simple introspection can often tell people that they do have concepts of things that they do not have words for. The American psychologist Greg Murphy relates that he regularly asks students in his courses which of them have a name for those clumps of dust that accumulate under beds on wooden floors. He typically finds that about half the class does (with *dust bunnies* and *dust monsters* being popular choices) while half the class does not. But the ones who do not have names for these things do recognize what Murphy is talking about, and so they presumably have DUST BUNNY concepts without corresponding words. Relatedly, new words, with previously unexpressed meanings, are regularly coined. (As I write, there is some concern about possible damage to internet services caused by *hacktivists*, who seem to be people who engage in social activism by means of computer hacking.) If our thought is constrained to run in channels carved out by our language, it is mysterious how people are able to formulate new concepts in this way.

As if these objections were not enough, there is excellent evidence of conceptual distinctions made by individuals who do not have any language. Susan Schaller, an American sign language interpreter and teacher, tells the story of meeting Ildefonso, a 27-year-old Mexican immigrant in Los Angeles who was congenitally deaf and had never acquired any language, spoken or signed. But he was clearly able to make conceptual distinctions: he had worked successfully in a number of agricultural jobs, for one thing. Eventually, after much patient work by Schaller, Ildefonso was able to become fluent in American Sign Language. And finally, psychologist Karen Wynn, in an experiment published in 1992, showed that babies as young as five months (well before the acquisition of number terminology) can do some elementary counting. Babies were shown a Mickey Mouse doll until they became bored and their eyes wandered. Then a screen came up and someone reached behind it with another doll and placed it there. When the screen came down, if there were two dolls present

(something the babies had not just seen) the babies looked for a few seconds and then were bored again: they plainly expected there to be two dolls behind the screen. But if there was just one doll (something that the babies had just been looking at) they were captivated: this evidently went counter to their expectations.

There is good reason to believe, then, that human beings can make conceptual distinctions that are not encoded by any languages they happen to speak.

What of the restricted Sapir-Whorf hypothesis? Recall that according to this view there are some topics such that the only conceptual distinctions we can make regarding them are those encoded in our language; and this comes about, as before, by our language imposing distinctions on the relevant sense data. This theory was also thought for a long time to be false. Recently, however, debate on it has been rekindled by some work on the recognition of numbers by the British psycholinguist Peter Gordon. In a 2004 article in the journal *Science*, Gordon reported that he had visited a remote hunter-gatherer tribe called the Pirahã, who live along the Maici River in the Lowland Amazonia region of Brazil. The Pirahã, of whom there are only about 250, are almost completely monolingual in their own language and have rejected assimilation into Brazilian society. Among other interesting properties, the Pirahã language has few or no number terms. The Pirahã words that look most like numbers are *hói* (with falling pitch), *hoí* (rising pitch) and *baagiso/aibaagi*. It is initially tempting to see these three words as forming a 'one-two-many' system of a type that has been documented in other languages: *hói* is often used for one object, *hoí* for two objects, and *baagiso* or *aibaagi* for more than that. But Gordon observed that 'whereas the word for "two" always denoted a larger quantity than the word for "one" (when used in the same context), the word for "one" was sometimes used to denote just a small quantity such as two or three or sometimes more.' A later study by Michael Frank, Daniel Everett, Evelina Fedorenko, and Edward Gibson found that *hói* was used for quantities of one to six, *hoí* for quantities of two to ten, and *baagiso* or *aibaagi* for quantities of three to ten. (They did not try to go above ten.) They

suggest that *hói* means something like 'few' rather than 'one' or even 'roughly one'.

It seems, then, that the Pirahã do not have words for numbers. They certainly do not have words for numbers above two. Gordon wanted to find out if, in spite of not having words for numbers above two, the Pirahã could nevertheless make distinctions between such numbers. He conducted several experiments to find out.

In one task, Gordon put some candy in a box with a small number of pictures of fish on the lid. There might be three fish on the lid, for example. He would then put the box behind his back and bring out two boxes: the original and another one identical except for the fact that it bore one fish more or less. He would ask his informants to pick out the box with the candy in it. When the task involved telling the difference between three and four fish, performance was just below 50%, i.e. around chance level. In another task, informants were asked to look at a small group of nuts for about eight seconds. Gordon then dropped the nuts into an opaque can into which the informants could not see. He withdrew them one at a time; and after each nut was taken out, informants were asked if there were any nuts left in the can. When two nuts had been placed in the can, the Pirahã correctly identified the moment when the can became empty about 90% of the time; when there were three nuts, they did so about 65% of the time; with four nuts, the figure was about 60%; with five, about 35%; and with six, about 10%.

So the answer Gordon arrived at was negative. The Pirahã could not reliably tell the difference between three and four, four and five, and so on. He concludes, 'The results of these studies show that the Pirahãs' impoverished counting system limits their ability to enumerate exact quantities when set sizes exceed two or three items.'

On the face of it, this is a classic Whorfian conclusion and strong support for the restricted Sapir-Whorf hypothesis. Number would be a topic such that the only conceptual distinctions we can make regarding it are those encoded in our language; since the Pirahã do not have any conceptual distinctions about number encoded in their

language, except words for 'few' and 'many', they cannot make any conceptual distinctions about number, except distinctions along the lines of 'few' and 'many'.

But it is worth asking how we can be sure about the direction of causation in this case. Remember that Whorf's idea, as expressed in my strong and restricted versions of the Sapir-Whorf hypothesis, was that our language imposes a set of distinctions on our thought. But what if our thought was constrained for some other reason? If there was some other reason why we could not (or would not) grasp a particular concept, it is natural to assume that our language would not have a word for that concept. Why would we have a word for something about which we had no idea? So it seems that the door is open for a sceptical alternative hypothesis: the Pirahã are unable (or unwilling) to think about numbers, and that is the reason their language does not have any words for numbers. The chain of causality goes from thought to language, in other words, and not vice versa.

This is certainly the line taken by the outside world's greatest expert on Pirahã language and culture, the American linguist Daniel Everett. Everett's own story is a remarkable one. He and his wife Keren first went to the Pirahã as Christian missionaries in 1977 and returned regularly at intervals for over twenty years. Previous attempts to evangelize this people had foundered due to the difficulty of their language and the lack of interest in outside cultures they displayed. The Everetts battled snakes, malaria, and occasional homicidal villagers in order to learn the Pirahã language and communicate God's word. Eventually Daniel Everett produced a translation of the Gospel of Mark into Pirahã and distributed recordings of it in hand-cranked cassette players. All to no avail. In fact, as Everett relates in his gripping book *Don't sleep, there are snakes*, the Pirahãs' pragmatic and sceptical attitude towards tales of old-time religion eventually played a large part in converting him: by the late 1980s he was an atheist.

Everett's take on the Pirahãs' inability to count is as follows. He claims that Pirahã culture is governed by a constraint called the *immediacy of experience principle*:

(2) Declarative Pirahã utterances contain only assertions directly related to the moment of speech, either experienced by the speaker or as witnessed by someone alive during the lifetime of the speaker.

This principle, if it obtains, would certainly seem to explain other interesting features of Pirahã language and culture noted by Everett: the fact that they have no creation myths and no attested stories of historical events that took place more than a generation or two ago. It is less clear, however, how it is supposed to explain the absence of numbers or the Pirahãs' inability to count. According to Everett, the connection is supposed to be that counting 'entails abstract generalizations that range in principle beyond immediate experience'. I think more would have to be said in order to make this explanation work. In principle, counting can of course involve abstract generalizations that range beyond immediate experience. ('Every Platonic Form is instantiated by twelve ancient megaliths on the dark side of the moon.') But it does not have to. A sincere utterance of 'Help, my left leg is being devoured by three panthers!' is rather directly related to the moment of speech as experienced by the speaker. So the immediacy of experience principle should not prohibit it, nor any one of an indefinitely large number of sentences involving numbers and immediate experience.

One could equally well argue about any word whatsoever that, since it could be used in an abstract generalization that ranges beyond immediate experience, it should therefore not be allowed in Pirahã. But Pirahã is not wholly devoid of words.

A more plausible alternative hypothesis has been advanced by the linguists Andew Nevins, David Pesetsky, and Cilene Rodrigues. They point out that there is a correlation between hunter-gatherer subsistence and restricted numeral systems. This makes sense, because such societies have no commerce or complex administration of a kind that would stimulate the invention of larger ones. So the Pirahã language has no numbers, according to this hypothesis, because the Pirahã people, given their culture, have no need of number concepts. Again, the direction of causality is reversed from that alleged by

Gordon. The Pirahã having no number concepts would, of course, explain their performance in Gordon's experiments.

It should be noted that this alternative hypothesis does not actually rule out the possibility that the direction of causality runs as Gordon alleged. We are simply left with two alternatives. But there is a certain naturalness in the alternative hypothesis that could well be thought to recommend it. So the restricted Sapir-Whorf hypothesis, at least as viewed in the light of Gordon's study, should be regarded as not proven.

What of the watered-down Sapir-Whorf hypothesis? Recall that this states that there are some topics such that the way we habitually or stereotypically think about them is influenced by the language we speak. This seems to be correct, although it should be emphasized that the degree to which we are influenced in this way may not be very great. Some of the best evidence in support of the watered-down Sapir-Whorf hypothesis comes from the work of the Russian psychologist Lera Boroditsky and her colleagues.

The experiments of Boroditsky's that I will review deal with grammatical gender. Languages that divide their nouns into masculine and feminine (and those that add neuter) frequently attribute masculine gender to nouns that name things that are biologically male and feminine to those that stand for the biological females. (Occasional exceptions are observed, of course.) Other nouns are divided up seemingly randomly, as Mark Twain observed in his wonderful essay 'The awful German language':

> [A] tree is male, its buds are female, its leaves are neuter; horses are sexless, dogs are male, cats are female,—Tom-cats included, of course... a person's nose, lips, shoulders, breast, hands, hips, and toes are of the female sex; and his hair, ears, eyes, chin, legs, knees, heart, and conscience, haven't any sex at all. The inventor of the language probably got what he knew about a conscience from hearsay.

Boroditsky wondered whether the masculine and feminine gender of the nouns that designated things that were not biologically male or female had any effect on how people thought about the things in question. Anecdotal evidence indicated that there might

be an effect here. Here is an extract from Andrei Makine's novel *Dreams of My Russian Summers*, as translated by Geoffrey Strachan. The protagonist is Russian, and his grandmother Charlotte is French:

> As a child I had absorbed all the sounds of Charlotte's language. I swam in them, without wondering why that glint in the grass, that colored, scented, living brilliance, sometimes existed in the masculine and had a crunchy, fragile, crystalline identity, imposed, it seemed, by one of its names, tsvetok; and was sometimes enveloped in a velvety, feltlike, and feminine aura, becoming une fleur.

The Russian word for 'flower', *tsvetok*, is masculine, whereas the French word, *fleur* is feminine.

Is it only sensitive young novelists learning French from their grandmothers who feel this kind of thing, or is the effect more general? Lera Boroditsky, Lauren Schmidt, and Webb Phillips, in a study published in 2003, performed two experiments to find out. Firstly, they investigated the effect of grammatical gender on memory. They used German and Spanish, two languages with grammatical gender which are such that many common object-names have one gender in German and a different one in Spanish, as with the example of 'flower' in French and Russian. For example, the word for 'apple' is masculine in German (*der Apfel*) and feminine in Spanish (*la manzana*).

The experimenters assembled twenty-four English nouns naming inanimate objects for which the corresponding nouns in Spanish and German had inconsistent genders. The subjects in the experiment were native speakers of Spanish and German. Both groups were competent in English, and the experiment was conducted in English. The experimenters taught each subject personal names for each object. For example an apple might be called Patrick or Patricia. For any given object, some speakers of German would be told a male name for it and some a female name; and likewise for the Spanish speakers. So some Spanish speakers had to memorize a male name like Patrick for the apple, while other Spanish speakers had to memorize a female name; and likewise for the German speakers.

The subjects were then tested on how well they remembered the personal names given to the objects. The result was that when a speaker had been given a name for an object that was of the same gender as the word for that kind of object in their native language, they were better at remembering the name. When the name was inconsistent in gender with the word in their native language, they were worse at remembering it than in the other case. For example, German speakers were better at remembering a male name like Patrick for an apple than they were at remembering a female name like Patricia. Conversely, Spanish speakers were better at remembering female names for apples than they were male names.

Why might this have happened? Boroditsky and her colleagues suggest that children learning a language that has grammatical gender have no reason to suppose that masculine and feminine genders of nouns do not reflect some real properties of the objects concerned, especially when they see biological males lining up with the masculine gender and biological females with the feminine. So, perhaps in an effort to remember the genders of the different nouns, children start focussing on stereotypically masculine qualities of objects named by masculine nouns and on stereotypically feminine qualities of objects named by feminine nouns. The result is that the conceptual representations that are the meanings of words (according to psychologists and Chomskyan linguists) are actually affected: the objects are portrayed, as it were, as having masculine or feminine properties. And according to contemporary psychology, there is no set of general-purpose conceptual representations that are somehow separate from the ones involved in word meaning. The concept APPLE that is the meaning of the word *apple* just is our concept of apples, with no back-ups or alternatives in sight. So learning a language with grammatical gender can affect your concepts. And with their concepts in this condition, it is small wonder that the Spanish-speaking and German-speaking subjects behaved as they did: they found it harder to remember names that disagreed in gender with the conceptual representations of the things they were names for.

In case you are wondering how, in concrete terms, one might go about portraying inanimate objects as masculine or feminine, here is an example: Boroditsky and her colleagues suggest that a toaster might be portrayed with an emphasis on its hard, metallic, and technological aspects in the minds of people for whom the corresponding word is masculine; but with an emphasis on its domestic and nourishing properties in the case of a feminine word.

Now this might be starting to sound like an awful lot of weight to put on one memory test, so Boroditsky and her colleagues tried to probe more directly to see what characteristics are associated with various objects by Spanish- and German-speakers. Using the list of 24 objects from the previous experiment, they asked Spanish- and German-speakers to write down the first three adjectives that came into their minds to describe each object on the list. After all the results were in, a group of English-speakers was asked to rate all the adjectives used according to whether they portrayed a masculine or a feminine quality. The results were that subjects produced adjectives that were rated more masculine for nouns that were masculine in their language, and more feminine for nouns that were feminine in their language.

Some of the examples were extremely striking. The word for 'key' is masculine in German (*der Schlüssel*) and feminine in Spanish (*la llave*). German-speakers used adjectives like *hard, heavy, jagged, metal, serrated,* and *useful.* Spanish-speakers used adjectives like *golden, intricate, little, lovely, shiny,* and *tiny.* The word for 'bridge' is feminine in German (*die Brücke*), masculine in Spanish (*el puente*). German-speakers used adjectives like *beautiful, elegant, fragile, peaceful, pretty,* and *slender.* Spanish-speakers (and you might like to imagine a deep manly voice at this point) used adjectives like *big, dangerous, long, strong, sturdy,* and *towering.* This, then, was a spectacular result.

Nevertheless, one might conceivably argue against the Whorfian force of these results by alleging that the direction of causation runs the other way, as with the Pirahã case. Perhaps some previously unremarked features of German and Hispanophone culture lead the

relevant people to regard bridges as having feminine and masculine traits, respectively, and the respective languages spoken in these cultures merely reflect this. It is entirely unclear what these cultural features could be; but to lay such doubts to rest, Boroditsky and her colleagues decided to invent a language that had a gender distinction and have people learn it. Since there was no culture attached to this particular language, culture could not be held responsible for any effects that were observed.

Boroditsky, Schmidt, and Phillips taught native English-speakers elements of a fictional language called Gumbuzi. The subjects were told that Gumbuzi nouns were divided into two classes called *oosative* and *soupative*. There were two words for 'the', *oos* or *sou*, and which one was used with which noun had to be memorized. (It can be seen that this is exactly how some languages with gender systems behave: compare German *der* vs *die* for 'the', and Spanish *el* vs *la*.) Nouns denoting male people and female people always clustered together, although the class they were said to belong to differed between subjects, so that one subject might learn that girls and women were oosative and men and boys were soupative, while another learnt the opposite. As with natural language gender systems, the categories extended to other objects too, so that one subject might learn that pans, forks, pencils, ballerinas, and girls were soupative while pots, spoons, pens, giants, and boys were oosative. Another, again, might learn the converse. The subjects were instructed in the distinction by means of labelled pictures on a computer screen (the picture of the giant clearly featuring a male giant) and were tested on twenty items until they could correctly tell for each one whether it was oosative or soupative.

When the subjects had learnt the distinction, they were shown the pictures again, without labels, and asked to come up with adjectives to describe them, as in the second experiment on Spanish- and German-speakers. As before, the adjectives were rated by an independent group as to whether they designated stereotypically male or stereotypically female characteristics.

As predicted, the subjects described inanimate objects that were grouped with biological females in the version of the gender system

they had learnt in terms that were more stereotypically feminine than the ones used for the objects that were grouped with the biological males. Once again, some of the differences were quite striking. When a violin was assigned feminine gender in the test, it was described as *artsy, beautiful, creative, curvy, delicate, elegant, interesting, pretty,* and *wooden.* When it was grouped with the biological males, it was described as *chirping, difficult, impressive, noisy, overused, piercing, shiny, slender, voluptuous,* and *wooden.* Someone's attention evidently wandered there in producing *slender* and *voluptuous*, but that is the price one pays for using human beings in one's experiments.

So even when a gender system is one that was made up specially for an experiment, with no accompanying culture in sight, it influences how people think of the things whose names are assigned to the various genders.

In other work, Boroditsky has demonstrated similar effects stemming from other parts of language. In a 2001 publication, for example, she demonstrated that Mandarin-speakers and English-speakers have different conceptions of time, arguably deriving from the spatial metaphors that are used about time in their respective languages. In English, we tend to use horizontal metaphors: next week's faculty board meeting is looming ominously *ahead* of us, whereas last week's is safely *behind* us, and so on. In Mandarin, although horizontal metaphors for time are common, it is also common to use vertical metaphors: earlier events are said to be *shàng* 'up', while later ones are said to be *xià* 'down'. Boroditsky showed native Mandarin- and English-speakers pictures showing spatial relationships and asked questions about them. For example, there might be a picture of a black ball above a white ball and the subjects would have to say whether or not the white ball was below the black one; and similarly for pictures showing horizontal spatial relationships. The subjects were then asked questions about time, like 'Does March come earlier than April?' The length of time the subjects took to answer was measured with a computer. The two groups responded very differently. The English-speakers responded more quickly to the temporal questions when they were preceded by a question about horizontal spatial relationships, whereas the

Mandarin-speakers were quicker when they had just answered a question about vertical spatial relationships. Boroditsky explained the difference in terms of priming, or the activation of certain mental representations, a term that you may remember from the description of Pylkkänen's experiment in Chapter 3: working out the answer to a question about horizontal or vertical spatial relationships made the subjects more ready to think in those terms; the fact that this had an effect on their ability to think about time shows that people's mental representations of time are the same as or similar to their mental representations of space, as the use of spatial metaphors for time suggests; and the fact that Mandarin-speakers were primed by vertical spatial relations and English-speakers by horizontal ones shows that the two groups represent time differently. Why should they do this? Well, one group has been brought up speaking a language that commonly uses vertical spatial metaphors for time, and the other group has been brought up with only horizontal spatial metaphors for time. Interestingly, a group of native English-speakers who were trained to speak about time vertically (with sentences like *Monday is above Tuesday*) produced results very like those of the Mandarin-speakers when they were tested immediately after training.

To what extent, then, does the language we speak influence the way we think? The strong Sapir-Whorf hypothesis is unequivocally false. The restricted Sapir-Whorf hypothesis is not proven. The only version of the hypothesis that has been shown to have something going for it is the watered-down one.

And here, although the results from Boroditsky's experiments are fascinating and impressive as far as they go, it is important to bear in mind that we are in general talking about exceedingly small differences that would not be noticed if they were not probed experimentally. Take the memory test featuring apples named Patrick. When the subjects were remembering a name that was consistent with the object's gender according to their native language, their success rate was 86%; when the name was inconsistent with the object's gender in their native language, their success rate was 78%. The difference was statistically significant, but not exactly indicative of a drastic loss of mental capacity. In the experiment on Mandarin-

and English-speakers' representations of time, the differences that showed up were a matter of milliseconds: Mandarin-speakers took an average of 2.347 seconds to answer the temporal questions after vertical priming, for example, and 2.503 seconds to answer them after horizontal priming, a difference of 156 milliseconds. Again, this was statistically significant in the context of the experiment, but it is not something that anyone would ever notice in real life.

The different rosters of adjectives provided by the Spanish- and German-speakers for keys and bridges constitute perhaps the most dramatic result we have surveyed. But here too it is important to bear in mind the limitations of the finding. At most, it shows that the ways in which these two groups of people habitually think about bridges (and so on) is influenced by their native language. But no-one is claiming that German-speakers, for example, are incapable of appreciating that bridges can be long, strong, sturdy, and towering. And, as cognitive scientist Steven Pinker has remarked in this context, 'just because a German thinks a bridge is feminine, doesn't mean he's going to ask one out on a date.'

The currently available evidence, then, suggests that the language we speak influences the way we think only with respect to scarcely perceptible cognitive biases that can be measured only in milliseconds and subtle stereotypes that vanish instantly upon reflection.

longior, paragraphs 128–32). The beating was introduced when the English philosopher Peter Geach turned the example into *Any man who owns a donkey beats it* on page 117 of his book *Reference and Generality* (Cornell University Press, 1962). This is still fairly old, though not nearly as old as the example itself. Geach was responsible for introducing these examples to modern philosophers and linguists. I owe most of this mini-history of donkey anaphora to Pieter Seuren's book *The Logic of Language* (Oxford University Press, 2010), page 300.

Descriptive indexicals were discovered by the American linguist Geoff Nunberg. See his article 'Indexicality and deixis', *Linguistics and Philosophy* 16: 1–43, 1993. The example about the Pope being an Italian, credited to Nunberg, actually comes from François Recanati's paper 'Deixis and anaphora', pages 286–316 of *Semantics versus Pragmatics* (Oxford University Press, 2005), edited by Zoltan Szabó. These technical works are aimed at professional linguists and philosophers.

The quotation from J.K Rowling is from page 279 of *Harry Potter and the Deathly Hallows* (Bloomsbury, 2007). The observation that pronouns can take on quantifier phrase meanings in this way has not been published before as far as I know.

The example about the Saturday before classes start is from page 29 of the article of Nunberg's just referred to.

Reams have been written about *It's raining*. As far as I know, the example originates in John Perry's article 'Thought without representation', *Supplementary Proceedings of the Aristotelian Society* 60 (1986): 263–283. The example *Everyone was sick* comes from Stephen Neale's book *Descriptions*, cited above in connection with Chapter 5, pages 94–5. And *I haven't eaten* is discussed on pages 135–6 of Kent Bach's article 'Conversational implicature', *Mind & Language* 9 (1994): 124–162. The theory about topic situations comes from Barwise and Perry's *Situations and Attitudes*, cited above in connection with Chapter 4, page 161. Influential objections to this theory along the lines given in the main text were put forward by Dag Westerståhl and Scott Soames: in Westerståhl's article 'Determiners and context sets', pages 45–71 of *Generalized Quantifiers in Natural Language* (Foris, 1985), edited by Johan van Benthem and Alice ter Meulen; and in Soames's article 'Incomplete definite descriptions', *Notre Dame Journal of Formal Logic* 27 (1986): 349–375. The example about the sleepers and the research assistant is from Soames. Covert indexicals in the syntax of the kind described here go back at least to Kai von Fintel's PhD dissertation *Restrictions on Quantifier Domains* (University of Massachusetts at Amherst, 1994). With the possible exception of

Neale's book, as noted above, none of this material is particularly accessible to beginners.

Jason Stanley's position can be appreciated by reading some or all of the essays in his book *Language in Context* (Oxford University Press, 2007). Sperber and Wilson's classic work is *Relevance: Communication and Cognition* (Blackwell, 1986). These two volumes would be fairly accessible to the reader of the current book, although the going would be rough in places.

Smith v. *United States* (91-8674), 508 U.S. 223 (1993). The case is briefly discussed on pages 23–4 of Antonin Scalia's book *A Matter of Interpretation: Federal Courts and the Law* (Princeton University Press, 1998). This is a fascinating and highly accessible guide to one justice's philosophy of legal interpretation. When I say that philosophers are starting to work on the issue of implicit content and legal interpretation, I am thinking chiefly of Stephen Neale, most of whose work on this subject is unfortunately unpublished. But there is a brief discussion of this case at the start of Neale's essay 'On location', pages 251–393 of *Situating Semantics: Essays on the Philosophy of John Perry* (MIT Press, 2007), edited by Michael O'Rourke and Corey Washington. This essay, which is mostly about *It's raining*, would be partly accessible to someone who has read the present book, but it would be rough going at times.

H.P. Grice's classic account of implicature is 'Logic and conversation', pages 41–58 of *Syntax and Semantics, 3: Speech Acts* (Academic Press, 1975), edited by P. Cole and J. Morgan. Some of this paper, as the title indicates, deals with the meanings of logical operators; but most of it would be quite accessible to the reader of the current book who has no more logical knowledge.

The full transcript of the speech of Hillary Clinton's quoted in this chapter can be found at http://transcripts.cnn.com/TRANSCRIPTS/0801/06/cnr.01.html.

Chapter 8

The Wittgenstein quotation is proposition 5.6 of the *Tractatus Logico-Philosophicus* in the translation of C.K. Ogden (Routledge & Kegan Paul, 1922).

The quotation from Whorf can be found on page 213 of *Language, Thought and Reality: Selected Writings of Benjamin Lee Whorf* (MIT Press, 1956), edited by John B. Carroll. Interestingly enough, Sapir used exactly the same adverb in one of his hedges on the subject: 'We see and hear and otherwise experience *very largely* as we do because the language habits of our

community predispose certain choices of interpretation' (my emphasis; page 210 of Sapir's article 'The status of linguistics as a science', *Language* 5(4), 1929, pages 207–14). Sapir and Whorf's writings on this subject are highly accessible.

Greg Murphy writes about dust bunnies on pages 389–90 of his book *The Big Book of Concepts*, cited above in connection with Chapter 1.

Susan Schaller's account of giving language to Ildefonso can be found in her lucid and accessible book *A Man Without Words* (Summit Books, 1991). Karen Wynn's experiment was published in her paper 'Addition and subtraction in human infants', *Nature* 358, 1992, pages 749–50. This paper is slightly technical, but should be mostly accessible to readers of the current book. I was alerted to the publications by Schaller and Wynn by reading the discussion of the Sapir-Whorf hypothesis in Chapter 3 of Steven Pinker's *The Language Instinct* (William Morrow, 1994), where I was also reminded of the relevance of Newspeak; in the unlikely event that you have not read this magnificently entertaining and persuasive book, I recommend that you do so soon.

The literature on Pirahã is already quite large and involves topics other than the number system. Amongst other highlights, Daniel Everett has alleged that Pirahã does not have recursion, the syntactic device that enables users of every other known human language to embed sentences inside other sentences (as subordinate clauses). But the literature is in general fascinating and more accessible than many technical publications. I estimate that the following works would be mostly accessible to the reader of the present book, although the going would be rough occasionally. Peter Gordon's article is 'Numerical cognition without words: evidence from Amazonia', *Science* 306, 2004, pages 496–99. The study by Michael Frank, Daniel Everett, Evelina Fedorenko, and Edward Gibson is 'Number as a cognitive technology: evidence from Pirahã language and cognition', *Cognition* 108, 2008, pages 819–24. Doubts about the direction of causation were first expressed in print by Daniel Casasanto in a letter to *Science*: 'Crying Whorf', *Science* 307, 2005, pages 1721–2 (followed by a reply by Gordon). Daniel Everett's book *Don't sleep, there are snakes* (Profile Books, 2008) is particularly engaging and accessible. (I disagree with some of what he has to say, though, as I mention in the main text.) Everett's main scholarly publication on Pirahã language and culture, including the immediacy of experience principle, is 'Cultural constraints on grammar and cognition in Pirahã: another look at the design features of human language', *Current Anthropology* 46(4), 2005, pages 621–34; it is followed by commentary by several other scholars and

a reply by Everett. The work cited by Nevins, Pesetsky, and Rodrigues, in spite of its relevance for Gordon's claims, is mainly devoted to criticizing Everett's work: 'Pirahã exceptionality: a reassessment', *Language* 85(2), 2009, pages 355–404. Everett replied in 'Pirahã culture and grammar: a response to some criticisms', *Language* 85(2), 2009, pages 405–42; the formulation of the immediacy of experience principle given in the text is taken from this article. Nevins, Pesetsky, and Rodrigues responded in turn with 'Evidence and argumentation: a reply to Everett (2009)', *Language* 85(3), 2009, pages 671–81.

I owe the quotation from Andrei Makine to a lecture by Lera Boroditsky before the Long Now Foundation given on 26 October 2010. It is available online at http://fora.tv/2010/10/26/Lera_Boroditsky_How_Language_Shapes_Thought. The paper by Boroditsky, Schmidt, and Phillips is 'Sex, syntax, and semantics', pages 61–79 of *Language in mind: advances in the study of language and thought* (MIT Press, 2003), edited by Dedre Gentner and Susan Goldin-Meadow. Boroditsky's article on Mandarin and English is 'Does language shape thought? Mandarin and English speakers' conceptions of time' in *Cognitive Psychology* 43: 1–22, 2001. The first of these publications is quite accessible; the second is rather more technical. The quotation from Steven Pinker is from a Boston Globe article from 18 November 2003, 'She explores the world of language and thought', available online at http://www-psych.stanford.edu/~lera/press/globe2003.html.