shapes on a video screen, because the user might later want to move the circle around and leave the triangle in place, or make the circle bigger or smaller, and one long list of dots would not allow the program to know which dots belong to the circle and which to the triangle. Instead, the shapes would be stored in some more abstract format (like the coordinates of a few defining points for each shape), a format that mirrors neither the inputs nor the outputs to the program but that can be translated to and from them when the need arises.

Grammar, a form of mental software, must have evolved under similar design specifications. Though psychologists under the influence of empiricism often suggest that grammar mirrors commands to the speech muscles, melodies in speech sounds, or mental scripts for the ways that people and things tend to interact, I think all these suggestions miss the mark. Grammar is a protocol that has to interconnect the ear, the mouth, and the mind, three very different kinds of machine. It cannot be tailored to any of them but must have an abstract logic of its own.

The idea that the human mind is designed to use abstract variables and data structures used to be, and in some circles still is, a shocking and revolutionary claim, because the structures have no direct counterpart in the child's experience. Some of the organization of grammar would have to be there from the start, part of the language-learning mechanism that allows children to make sense out of the noises they hear from their parents. The details of syntax have figured prominently in the history of psychology, because they are a case where complexity in the mind is not caused by learning; learning is caused by complexity in the mind. And that was real news.

5

Words, Words, Words

The word glamour *comes from the word* grammar, *and since the* Chomskyan revolution the etymology has been fitting. Who could not be dazzled by the creative power of the mental grammar, by its ability to convey an infinite number of thoughts with a finite set of rules? There has been a book on mind and matter called *Grammatical Man,* and a Nobel Prize lecture comparing the machinery of life to a generative grammar. Chomsky has been interviewed in *Rolling Stone* and alluded to on *Saturday Night Live.* In Woody Allen's story "The Whore of Mensa," the patron asks, "Suppose I wanted Noam Chomsky explained to me by two girls?" "It'd cost you," she replies.

Unlike the mental grammar, the mental dictionary has had no cachet. It seems like nothing more than a humdrum list of words, each transcribed into the head by dull-witted rote memorization. In the preface to his *Dictionary,* Samuel Johnson wrote:

> It is the fate of those who dwell at the lower employments of life, to be rather driven by the fear of evil, than attracted by the prospect of good; to be exposed to censure, without hope of praise; to be disgraced by miscarriage, or punished for neglect, where success would have been without applause, and diligence without reward.

> Among these unhappy mortals is the writer of dictionaries.

Johnson's own dictionary defines *lexicographer* as "a harmless drudge, that busies himself in tracing the original, and detailing the signification of words."

In this chapter we will see that the stereotype is unfair. The world of words is just as wondrous as the world of syntax, or even more so. For not only are people as infinitely creative with words as they are with phrases and sentences, but memorizing individual words demands its own special virtuosity.

Recall the *wug*-test, passed by any preschooler: "Here is a wug. Now there are two of them. There are two___." Before being so challenged, the child has neither heard anyone say, nor been rewarded for saying, the word *wugs*. Therefore words are not simply retrieved from a mental archive. People must have a mental rule for generating new words from old ones, something like "To form the plural of a noun, add the suffix *-s*." The engineering trick behind human language—its being a discrete combinatorial system—is used in at least two different places: sentences and phrases are built out of words by the rules of syntax, and the words themselves are built out of smaller bits by another set of rules, the rules of "morphology."

The creative powers of English morphology are pathetic compared to what we find in other languages. The English noun comes in exactly two forms *(**duck** and **ducks)*, the verb in four *(**quack, quacks, quacked, quacking**)*. In modern Italian and Spanish every verb has about fifty forms; in classical Greek, three hundred and fifty; in Turkish, two million! Many of the languages I have brought up, such as Eskimo, Apache, Hopi, Kivunjo, and American Sign Language, are known for this prodigious ability. How do they do it? Here is an example from Kivunjo, the Bantu language that was said to make English look like checkers compared to chess. The verb "Näïkì-ńlyïïà," meaning "He is eating it for her," is composed of eight parts:

- N-: A marker indicating that the word is the "focus" of that point in the conversation.

- -ä-: A subject agreement marker. It identifies the eater as falling into Class 1 of the sixteen gender classes, "human singular." (Remember that to a linguist "gender" means kind, not sex.) Other genders embrace nouns that pertain to several humans, thin or extended objects, objects that come in pairs or clusters, the pairs or clusters themselves, instruments, animals, body parts, diminutives (small or cute versions of things), abstract qualities, precise locations, and general locales.
- -ï-: Present tense. Other tenses in Bantu can refer to today, earlier today, yesterday, no earlier than yesterday, yesterday or earlier, in the remote past, habitually, ongoing, consecutively, hypothetically, in the future, at an indeterminate time, not yet, and sometimes.
- -kì-: An object agreement marker, in this case indicating that the thing eaten falls into gender Class 7.
- -ń-: A benefactive marker, indicating for whose benefit the action is taking place, in this case a member of gender Class 1.
- -lyì-: The verb, "to eat."
- -ï-: An "applicative" marker, indicating that the verb's cast of players has been augmented by one additional role, in this case the benefactive. (As an analogy, imagine that in English we had to add a suffix to the verb *bake* when it is used in *I **baked her a cake*** as opposed to the usual *I **baked a cake**.*)
- -à: A final vowel, which can indicate indicative versus subjunctive mood.

If you multiply out the number of possible combinations of the seven prefixes and suffixes, the product is about half a million, and that is the number of possible forms per verb in the language. In effect, Kivunjo and languages like it are building an entire sentence inside a single complex word, the verb.

But I have been a bit unfair to English. English is genuinely

crude in its "inflectional" morphology, where one modifies a word to fit the sentence, like marking a noun for the plural with *-s* or a verb for past tense with *-ed*. But English holds its own in "derivational" morphology, where one creates a new word out of an old one. For example, the suffix *-able*, as in *learnable, teachable,* and *huggable,* converts a verb meaning "to do X" into an adjective meaning "capable of having X done to it." Most people are surprised to learn how many derivational suffixes there are in English. Here are the more common ones:

-able	-ate	-ify	-ize
-age	-ed	-ion	-ly
-al	-en	-ish	-ment
-an	-er	-ism	-ness
-ant	-ful	-ist	-ory
-ance	-hood	-ity	-ous
-ary	-ic	-ive	-y

In addition, English is free and easy with "compounding," which glues two words together to form a new one, like *toothbrush* and *mouse-eater*. Thanks to these processes, the number of possible words, even in morphologically impoverished English, is immense. The computational linguist Richard Sproat compiled all the distinct words used in the forty-four million words of text from Associated Press news stories beginning in mid-February 1988. Up through December 30, the list contained three hundred thousand distinct word forms, about as many as in a good unabridged dictionary. You might guess that this would exhaust the English words that would ever appear in such stories. But when Sproat looked at what came over the wire on December 31, he found no fewer than thirty-five new forms, including *instrumenting, counterprograms, armhole, part-Vulcan, fuzzier, grooveled, boulderlike, mega-lizard, traumatological,* and *ex-critters.*

Even more impressive, the output of one morphological rule can be the input to another, or to itself: one can talk about the *unmicrowaveability* of some French fries or a *toothbrush-holder fastener box* in which to keep one's toothbrush-holder fasteners. This makes the

number of possible words in a language bigger than immense; like the number of sentences, it is infinite. Putting aside fanciful coinages concocted for immortality in *Guinness,* a candidate for the longest word to date in English might be *floccinaucinihilipilification,* defined in the *Oxford English Dictionary* as "the categorizing of something as worthless or trivial." But that is a record meant to be broken:

> *floccinaucinihilipilificational:* pertaining to the categorizing of something as worthless or trivial
>
> *floccinaucinihilipilificationalize:* to cause something to pertain to the categorizing of something as worthless or trivial
>
> *floccinaucinihilipilificationalization:* the act of causing something to pertain to the categorizing of something as worthless or trivial
>
> *floccinaucinihilipilificationalizational:* pertaining to the act of causing something to pertain to the categorizing of something as worthless or trivial
>
> *floccinaucinihilipilificationalizationalize:* to cause something to pertain to the act of causing something to pertain . . .

Or, if you suffer from sesquipedaliaphobia, you can think of your *great-grandmother,* your *great-great-grandmother,* your *great-great-great-grandmother,* and so on, limited only in practice by the number of generations since Eve.

What's more, words, like sentences, are too delicately layered to be generated by a chaining device (a system that selects an item from one list, then moves on to some other list, then to another). When Ronald Reagan proposed the Strategic Defense Initiative, popularly known as Star Wars, he imagined a future in which an incoming Soviet missile would be shot down by an *anti-missile missile.* But critics pointed out that the Soviet Union could counterattack with an *anti-anti-missile-missile missile.* No problem, said his MIT-educated engineers; we'll just build an *anti-anti-anti-missile-missile-missile missile.* These high-tech weapons need a high-tech grammar—something that can keep track of all the *anti*'s at the beginning of the word so that it can complete the word with an equal number of *missile*'s, plus one, at

the end. A word structure grammar (a phrase structure grammar for words) that can embed a word in between an *anti-* and its *missile* can achieve these objectives; a chaining device cannot, because it has forgotten the pieces that it laid down at the beginning of the long word by the time it gets to the end.

Like syntax, morphology is a cleverly designed system, and many of the seeming oddities of words are predictable products of its internal logic. Words have a delicate anatomy consisting of pieces, called morphemes, that fit together in certain ways. The word structure system is an extension of the X-bar phase structure system, in which big nounish things are built out of smaller nounish things, smaller nounish things are built out of still smaller nounish things, and so on. The biggest phrase involving nouns is the noun phrase; a noun phrase contains an N-bar; an N-bar contains a noun—the word. Jumping from syntax to morphology, we simply continue the dissection, analyzing the word into smaller and smaller nounish pieces.

Here is a picture of the structure of the word *dogs:*

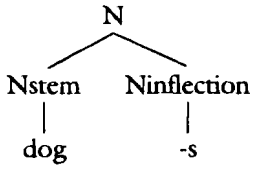

The top of this mini-tree is "N" for "noun"; this allows the docking maneuver in which the whole word can be plugged into the noun slot inside any noun phrase. Down inside the word, we have two parts: the bare word form *dog,* usually called the stem, and the plural inflection *-s.* The rule responsible for inflected words (the rule of *wug*-test fame) is simply

N → Nstem Ninflection
"A noun can consist of a noun stem followed by a noun inflection."

The rule nicely interfaces with the mental dictionary: *dog* would be listed as a noun stem meaning "dog," and *-s* would be listed as a noun inflection meaning "plural of."

This rule is the simplest, most stripped-down example of anything we would want to call a rule of grammar. In my laboratory we use it as an easily studied instance of mental grammar, allowing us to document in great detail the psychology of linguistic rules from infancy to old age in both normal and neurologically impaired people, in much the same way that biologists focus on the fruit fly *Drosophila* to study the machinery of genes. Though simple, the rule that glues an inflection to a stem is a surprisingly powerful computational operation. That is because it recognizes an abstract mental symbol, like "noun stem," instead of being associated with a particular list of words or a particular list of sounds or a particular list of meanings. We can use the rule to inflect any item in the mental dictionary that lists "noun stem" in its entry, without caring what the word means; we can convert not only *dog* to *dogs* but also *hour* to *hours* and *justification* to *justifications.* Likewise, the rule allows us to form plurals without caring what the word sounds like; we pluralize unusual-sounding words as in *the Gorbachevs, the Bachs,* and *the Mao Zedongs.* For the same reason, the rule is perfectly happy applying to brand-new nouns, like *faxes, dweebs, wugs,* and *zots.*

We apply the rule so effortlessly that perhaps the only way I can drum up some admiration for what it accomplishes is to compare humans with a certain kind of computer program that many computer scientists tout as the wave of the future. These programs, called "artificial neural networks," do not apply a rule like the one I have just shown you. An artificial neural network works by analogy, converting *wug* to *wugged* because it is vaguely similar to *hug–hugged, walk–walked,* and thousands of other verbs the network has been trained to recognize. But when the network is faced with a new verb that is unlike anything it has previously been trained on, it often mangles it, because the network does not have an abstract, all-embracing category "verb stem" to fall back on and add an affix to. Here are some com-

parisons between what people typically do and what artificial neural networks typically do when given a *wug*-test:

VERB	TYPICAL PAST-TENSE FORM GIVEN BY PEOPLE	TYPICAL PAST-TENSE FORM GIVEN BY NEURAL NETWORKS
mail	mailed	membled
conflict	conflicted	conflafted
wink	winked	wok
quiver	quivered	quess
satisfy	satisfied	sedderded
smairf	smairfed	sprurice
trilb	trilbed	treelilt
smeej	smeejed	leefloag
frilg	frilged	freezled

Stems can be built out of parts, too, in a second, deeper level of word assembly. In compounds like *Yugoslavia report, sushi-lover, broccoli-green,* and *toothbrush,*

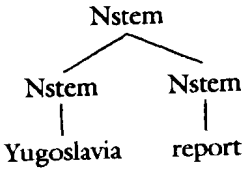

two stems are joined together to form a new stem, by the rule

Nstem → Nstem Nstem
"A noun stem can consist of a noun stem followed by another noun stem."

In English, a compound is often spelled with a hyphen or by running its two words together, but it can also be spelled with a space between the two components as if they were still separate words. This confused your grammar teacher into telling you that in *Yugoslavia report*, "Yugoslavia" is an adjective. To see that this can't be right, just try comparing it with a real adjective like *interesting*. You can say *This report seems interesting* but not *This report seems Yugoslavia*! There is a simple way to tell whether something is a compound word or a phrase: compounds generally have stress on the first word, phrases on the second. A *dark róom* (phrase) is any room that is dark, but a *dárk room* (compound word) is where photographers work, and a dark-room can be lit when the photographer is done. A *black bóard* (phrase) is necessarily a board that is black, but some *blackboards* (compound word) are green or even white. Without pronunciation or punctuation as a guide, some word strings can be read either as a phrase or as a compound, like the following headlines:

Squad Helps Dog Bite Victim
Man Eating Piranha Mistakenly Sold as Pet Fish
Juvenile Court to Try Shooting Defendant

New stems can also be formed out of old ones by adding affixes (prefixes and suffixes), like the *-al, -ize,* and *-ation* I used recursively to get longer and longer words ad infinitum (as in *sensationalizationalization*). For example, *-able* combines with any verb to create an adjective, as in *crunch–crunchable*. The suffix *-er* converts any verb to a noun, as in *crunch–cruncher,* and the suffix *-ness* converts any adjective into a noun, as in *crunchy–crunchiness*.

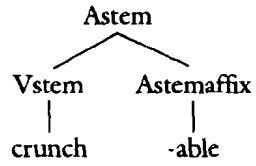

The rule forming them is

Astem → Stem Astemaffix
"An adjective stem can consist of a stem joined to a suffix."

and a suffix like *-able* would have a mental dictionary entry like the following:

-able:
adjective stem affix
means "capable of being X'd"
attach me to a verb stem

Like inflections, stem affixes are promiscuous, mating with any stem that has the right category label, and so we have *crunchable, scrunchable, shmooshable, wuggable,* and so on. Their meanings are predictable: capable of being crunched, capable of being scrunched, capable of being shmooshed, even capable of being "wugged," whatever *wug* means. (Though I can think of an exception: in the sentence *I asked him what he thought of my review in his book, and his response was unprintable,* the word *unprintable* means something much more specific than "incapable of being printed.")

The scheme for computing the meaning of a stem out of the meaning of its parts is similar to the one used in syntax: one special element is the "head," and it determines what the conglomeration refers to. Just as the phrase *the cat in the hat* is a kind of cat, showing that *cat* is its head, a *Yugoslavia report* is a kind of report, and *shmooshability* is a kind of ability, so *report* and *-ability* must be the heads of those words. The head of an English word is simply its rightmost morpheme.

Continuing the dissection we can tease stems into even smaller parts. The smallest part of a word, the part that cannot be cut up into any smaller parts, is called its root. Roots can combine with special suffixes to form stems. For example, the root *Darwin* can be found inside the stem *Darwinian.* The stem *Darwinian* in turn can be fed into the suffixing rule to yield the new stem *Darwinianism.* From there, the inflectional rule could even give us the word *Darwinianisms,* embodying all three levels of word structure:

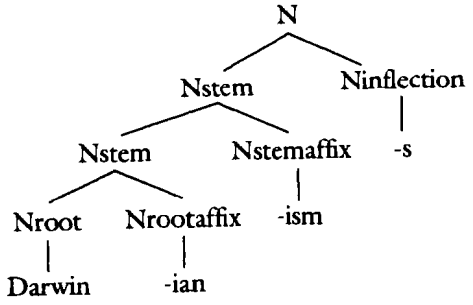

Interestingly, the pieces fit together in only certain ways. Thus *Darwinism,* a stem formed by the stem suffix *-ism,* cannot be a host for *-ian,* because *-ian* attaches only to roots; hence *Darwinismian* (which would mean "pertaining to Darwinism") sounds ridiculous. Similarly, *Darwinsian* ("pertaining to the two famous Darwins, Charles and Erasmus"), *Darwinsianism,* and *Darwinsism* are quite impossible, because whole inflected words cannot have any root or stem suffixes joined to them.

Down at the bottommost level of roots and root affixes, we have entered a strange world. Take *electricity.* It seems to contain two parts, *electric* and *-ity:*

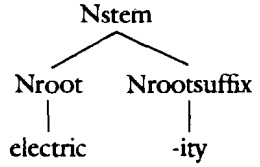

But are these words really assembled by a rule, gluing a dictionary entry for *-ity* onto the root *electric,* like this?

Nstem → Nroot Nrootsuffix
"A noun stem can be composed of a noun root and a suffix."

-ity:
noun root suffix
means "the state of being X"
attach me to a noun root

Not this time. First, you can't get *electricity* simply by gluing together the word *electric* and the suffix *-ity*—that would sound like "electrick itty." The root that *-ity* is attached to has changed its pronunciation to "electríss." That residue, left behind when the suffix has been removed, is a root that cannot be pronounced in isolation.

Second, root-affix combinations have unpredictable meanings; the neat scheme for interpreting the meaning of the whole from the meaning of the parts breaks down. *Complexity* is the state of being complex, but *electricity* is not the state of being electric (you would

never say that the electricity of this new can opener makes it convenient); it is the force powering something electric. Similarly, *instrumental* has nothing to do with instruments, *intoxicate* is not about toxic substances, one does not recite at a *recital*, and a five-speed *transmission* is not an act of transmitting.

Third, the supposed rule and affix do not apply to words freely, unlike the other rules and affixes we have looked at. For example, something can be *academic* or *acrobatic* or *aerodynamic* or *alcoholic*, but *academicity, acrobaticity, aerodynamicity,* and *alcoholicity* sound horrible (to pick just the first four words ending in *-ic* in my electronic dictionary).

So at the third and most microscopic level of word structure, roots and their affixes, we do not find bona fide rules that build words according to predictable formulas, *wug*-style. The stems seem to be stored in the mental dictionary with their own idiosyncratic meanings attached. Many of these complex stems originally were formed after the Renaissance, when scholars imported many words and suffixes into English from Latin and French, using some of the rules appropriate to those languages of learning. We have inherited the words, but not the rules. The reason to think that modern English speakers mentally analyze these words as trees at all, rather than as homogeneous strings of sound, is that we all sense that there is a natural break point between the *electric* and the *-ity*. We also recognize that there is an affinity between the word *electric* and the word *electricity*, and we recognize that any other word containing *-ity* must be a noun.

Our ability to appreciate a pattern inside a word, while knowing that the pattern is not the product of some potent rule, is the inspiration for a whole genre of wordplay. Self-conscious writers and speakers often extend Latinate root suffixes to new forms by analogy, such as *religiosity, criticality, systematicity, randomicity, insipidify, calumniate, conciliate, stereotypy, disaffiliate, gallonage,* and *Shavian*. The words have an air of heaviosity and seriosity about them, making the style an easy target for parody. A 1982 editorial cartoon by Jeff MacNelly put the following resignation speech into the mouth of Alexander Haig, the malaprop-prone Secretary of State:

> I decisioned the necessifaction of the resignatory action/option due to the dangerosity of the trendflowing of foreign policy away from our originatious careful coursing towards consistensivity, purposity, steadfastnitude, and above all, clarity.

Another cartoon, by Tom Toles, showed a bearded academician explaining the reason verbal Scholastic Aptitude Test scores were at an all-time low:

> Incomplete implementation of strategized programmatics designated to maximize acquisition of awareness and utilization of communications skills pursuant to standardized review and assessment of languaginal development.

In the culture of computer programmers and managers, this analogy-making is used for playful precision, not pomposity. *The New Hacker's Dictionary,* a compilation of hackish jargon, is a near-exhaustive catalogue of the not-quite-freely-extendible root affixes in English:

> *ambimoustrous* adj. Capable of operating a mouse with either hand.
> *barfulous* adj. Something that would make anyone barf.
> *bogosity* n. The degree to which something is bogus.
> *bogotify* v. To render something bogus.
> *bozotic* adj. Having the quality of Bozo the Clown.
> *cuspy* adj. Functionally elegant.
> *depeditate* v. To cut the feet off of (e.g., while printing the bottom of a page).
> *dimwittery* n. Example of a dim-witted statement.
> *geekdom* n. State of being a techno-nerd.
> *marketroid* n. Member of a company's marketing department.
> *mumblage* n. The topic of one's mumbling.
> *pessimal* adj. Opposite of "optimal."
> *wedgitude* n. The state of being wedged (stuck; incapable of proceeding without help).
> *wizardly* adj. Pertaining to expert programmers.

Down at the level of word roots, we also find messy patterns in irregular plurals like *mouse–mice* and *man–men* and in irregular past-tense forms like *drink–drank* and *seek–sought*. Irregular forms tend to come in families, like *drink–drank, sink–sank, shrink–shrank, stink–stank, sing–sang, ring–rang, spring–sprang, swim–swam,* and *sit–sat,* or *blow–blew, know–knew, grow–grew, throw–threw, fly–flew,* and *slay–slew*. This is because thousands of years ago Proto-Indo-European, the language ancestral to English and most other European languages, had rules that replaced one vowel with another to form the past tense, just as we now have a rule that adds *-ed*. The irregular or "strong" verbs in modern English are mere fossils of these rules; the rules themselves are dead and gone. Most verbs that would seem eligible to belong to the irregular families are arbitrarily excluded, as we see in the following doggerel:

> Sally Salter, she was a young teacher who taught,
> And her friend, Charley Church, was a preacher who praught;
> Though his enemies called him a screecher, who scraught.
>
> His heart, when he saw her, kept sinking, and sunk;
> And his eye, meeting hers, began winking, and wunk;
> While she in her turn, fell to thinking, and thunk.
>
> In secret he wanted to speak, and he spoke,
> To seek with his lips what his heart long had soke,
> So he managed to let the truth leak, and it loke.
>
> The kiss he was dying to steal, then he stole;
> At the feet where he wanted to kneel, then he knole;
> And he said, "I feel better than ever I fole."

People must simply be memorizing each past-tense form separately. But as this poem shows, they can be sensitive to the patterns among them and can even extend the patterns to new words for humorous effect, as in Haigspeak and hackspeak. Many of us have been tempted by the cuteness of *sneeze–snoze, squeeze–squoze, take–took–tooken,* and *shit-shat,* which are based on analogies with *freeze–froze, break–broke–broken,* and *sit–sat*. In *Crazy English* Richard

Lederer wrote an essay called "Foxen in the Henhice," featuring irregular plurals gone mad: *booth–beeth, harmonica–harmonicae, mother–methren, drum–dra, Kleenex–Kleenices,* and *bathtub–bathtubim*. Hackers speak of *faxen, VAXen, boxen, meece,* and *Macinteesh*. *Newsweek* magazine once referred to the white-caped, rhinestone-studded Las Vegas entertainers as *Elvii*. In the *Peanuts* comic strip, Linus's teacher Miss Othmar once had the class glue eggshells into model *igli*. Maggie Sullivan wrote an article in the *New York Times* calling for "strengthening" the English language by conjugating more verbs as if they were strong:

> *Subdue, subdid, subdone:* Nothing could have subdone him the way her violet eyes subdid him.
>
> *Seesaw, sawsaw, seensaw:* While the children sawsaw, the old man thought of long ago when he had seensaw.
>
> *Pay, pew, pain:* He had pain for not choosing a wife more carefully.
>
> *Ensnare, ensnore, ensnorn:* In the 60's and 70's, Sominex ads ensnore many who had never been ensnorn by ads before.
>
> *Commemoreat, commemorate, commemoreaten:* At the banquet to commemoreat Herbert Hoover, spirits were high, and by the end of the evening many other Republicans had been commemoreaten.

In Boston there is an old joke about a woman who landed at Logan Airport and asked the taxi driver, "Can you take me someplace where I can get scrod?" He replied, "Gee, that's the first time I've heard it in the pluperfect subjunctive."

Occasionally a playful or cool-sounding form will catch on and spread through the language community, as *catch–caught* did several hundred years ago on the analogy of *teach–taught* and as *sneak–snuck* is doing today on the analogy of *stick–stuck*. (I am told that *has tooken* is the preferred form among today's mall rats.) This process can be seen clearly when we compare dialects, which retain the products of their own earlier fads. The curmudgeonly columist H. L. Mencken was also a respectable amateur linguist, and he documented many

past-tense forms found in American regional dialects, like *heat–het* (similar to *bleed–bled*), *drag–drug* (*dig–dug*), and *help–holp* (*tell–told*). Dizzy Dean, the St. Louis Cardinals pitcher and CBS announcer, was notorious for saying "He slood into second base," common in his native Arkansas. For four decades English teachers across the nation engaged in a letter-writing campaign to CBS demanding that he be removed, much to his delight. One of his replies, during the Great Depression, was "A lot of folks that ain't sayin' 'ain't' ain't eatin'." Once he baited them with the following play-by-play:

> The pitcher wound up and flang the ball at the batter. The batter swang and missed. The pitcher flang the ball again and this time the batter connected. He hit a high fly right to the center fielder. The center fielder was all set to catch the ball, but at the last minute his eyes were blound by the sun and he dropped it!

But successful adoptions of such creative extensions are rare; irregulars remain mostly as isolated oddballs.

Irregularity in grammar seems like the epitome of human eccentricity and quirkiness. Irregular forms are explicitly abolished in "rationally designed" languages like Esperanto, Orwell's Newspeak, and Planetary League Auxiliary Speech in Robert Heinlein's science fiction novel *Time for the Stars*. Perhaps in defiance of such regimentation, a woman in search of a nonconformist soulmate recently wrote this personal ad in the *New York Review of Books:*

> **Are you an irregular verb** who believes nouns have more power than adjectives? Unpretentious, professional DWF, 5 yr. European resident, sometime violinist, slim, attractive, with married children. . . . Seeking sensitive, sanguine, youthful man, mid 50's–60's, health-conscious, intellectually adventurous, who values truth, loyalty, and openness.

A general statement of irregularity and the human condition comes from the novelist Marguerite Yourcenar: "Grammar, with its mixture of logical rule and arbitrary usage, proposes to a young mind a foretaste of what will be offered to him later on by law and ethics, those sciences of human conduct, and by all the systems wherein man has codified his instinctive experience."

For all its symbolism about the freewheeling human spirit, though, irregularlity is tightly encapsulated in the word-building system; the system as a whole is quite cuspy. Irregular forms are roots, which are found inside stems, which are found inside words, some of which can be formed by regular inflection. This layering not only predicts many of the possible and impossible words of English (for example, why *Darwinianism* sounds better than *Darwinismian*); it provides a neat explanation for many trivia questions about seemingly illogical usage, such as: Why in baseball is a batter said to have *flied out*—why has no mere mortal ever *flown out* to center field? Why is the hockey team in Toronto called the *Maple Leafs* and not the *Maple Leaves?* Why do many people say *Walkmans,* rather than *Walkmen,* as the plural of *Walkman?* Why would it sound odd for someone to say that all of his daughter's friends are *low-lives?*

Consult any style manual or how-to book on grammar, and it will give one or two explanations as to why the irregular is tossed aside—both wrong. One is that the books are closed on irregular words in English; any new form added to the language must be regular. Not true: if I coin new words like *to re-sing* or *to out-sing,* their pasts are *re-sang* and *out-sang,* not *re-singed* and *out-singed.* Similarly, I recently read that there are peasants who run around with small tanks in China's oil fields, scavenging oil from unguarded wells; the article calls them *oil-mice,* not *oil-mouses.* The second explanation is that when a word acquires a new, nonliteral sense, like baseball's *fly out,* that sense requires a regular form. The oil-mice clearly falsify that explanation, as do the many other metaphors based on irregular

nouns, which steadfastly keep their irregularity: *sawteeth* (not *saw-tooths*), *Freud's intellectual children* (not *childs*), *snowmen* (not *snow-mans*), and so on. Likewise, when the verb *to blow* developed slang meanings like *to blow him away* (assassinate) and *to blow it off* (dismiss casually), the past-tense forms remained irregular: *blew him away* and *blew off the exam,* not *blowed him away* and *blowed off the exam.*

The real rationale for *flied out* and *Walkmans* comes from the algorithm for interpreting the meanings of complex words from the meanings of the simple words they are built out of. Recall that when a big word is built out of smaller words, the big word gets all its properties from one special word sitting inside it at the extreme right: the head. The head of the verb *to overshoot* is the verb *to shoot,* so *overshooting* is a kind of *shooting,* and it is a verb, because *shoot* is a verb. Similarly, a *workman* is a singular noun, because *man,* its head, is a singular noun, and it refers to a kind of man, not a kind of work. Here is what the word structures look like:

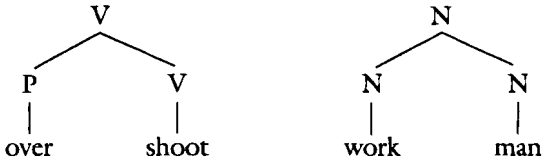

Crucially, the percolation conduit from the head to the top node applies to *all* the information stored with the head word: not just its nounhood or verbhood, and not just its meaning, but any irregular form that is stored with it, too. For example, part of the mental dictionary entry for *shoot* would say "I have my own irregular past-tense form, *shot.*" This bit of information percolates up and applies to the complex word, just like any other piece of information. The past tense of *overshoot* is thus *overshot* (not *overshooted*). Likewise, the word *man* bears the tag "My plural is *men.*" Since *man* is the head of *workman,* the tag percolates up to the N symbol standing for *workman,* and so the plural of *workman* is *workmen.* This is also why we get *out-sang, oil-mice, sawteeth,* and *blew him away.*

Now we can answer the trivia questions. The source of quirkiness in words like *fly out* and *Walkmans* is their *headlessness.* A headless word is an exceptional item that, for one reason or another, differs in some property from its rightmost element, the one it would be based on if it were like ordinary words. A simple example of a headless word is a *low-life*—not a kind of life at all but a kind of person, namely one who leads a low life. In the word *low-life,* then, the normal percolation pipeline must be blocked. Now, a pipeline inside a word cannot be blocked for just one kind of information; if it is blocked for one thing, nothing passes through. If *low-life* does not get its meaning from *life,* it cannot get its plural from *life* either. The irregular form associated with *life,* namely *lives,* is trapped in the dictionary, with no way to bubble up to the whole word *low-life.* The all-purpose regular rule, "Add the *-s* suffix," steps in by default, and we get *low-lifes.* By similar unconscious reasoning, speakers arrive at *saber-tooths* (a kind of tiger, not a kind of tooth), *tenderfoots* (novice cub scouts, who are not a kind of foot but a kind of youngster that has tender feet), *flatfoots* (also not a kind of foot but a slang term for policemen), and *still lifes* (not a kind of life but a kind of painting).

Since the Sony Walkman was introduced, no one has been sure whether two of them should be *Walkmen* or *Walkmans.* (The nonsexist alternative *Walkperson* would leave us on the hook, because we would be faced with a choice between *Walkpersons* and *Walkpeople.*) The temptation to say *Walkmans* comes from the word's being headless: a Walkman is not a kind of man, so it must not be getting its meaning from the word *man* inside it, and by the logic of headlessness it shouldn't receive a plural form from *man,* either. But it is hard to be comfortable with any kind of plural, because the relation between *Walkman* and *man* feels utterly obscure. It feels obscure because the word was not put together by any recognizable scheme. It is an example of the pseudo-English that is popular in Japan in signs and product names. (For example, one popular soft drink is called Sweat, and T-shirts have enigmatic inscriptions like CIRCUIT BEAVER, NURSE MENTALITY, and BONERACTIVE WEAR.) The Sony Corporation has an official answer to the question of how to refer to more than one Walkman. Fearing that their trademark, if converted to a noun, may

become as generic as *aspirin* or *kleenex*, they sidestep the grammatical issues by insisting upon *Walkman Personal Stereos.*

What about flying out? To the baseball cognoscenti, it is not directly based on the familiar verb *to fly* ("to proceed through the air") but on the noun *a fly* ("a ball hit on a conspicuously parabolic trajectory"). To *fly out* means "to make an out by hitting a fly that gets caught." The noun *a fly*, of course, itself came from the verb *to fly*. The word-within-a-word-within-a-word structure can be seen in this bamboo-like tree:

Since the whole word, represented by its topmost label, is a verb, but the element it is made out of one level down is a noun, *to fly out*, like *low-life*, must be headless—if the noun *fly* were its head, *fly out* would have to be a noun, too, which it is not. Lacking a head and its associated data pipeline, the irregular forms of the original verb *to fly*, namely *flew* and *flown*, are trapped at the bottommost level and cannot bubble up to attach to the whole word. The regular *-ed* rule rushes in in its usual role as the last resort, and thus we say that Wade Boggs *flied out*. What kills the irregularity of *to fly out*, then, is not its specialized meaning, but its being a verb based on a word that is not a verb. By the same logic, we say *They ringed the city with artillery* ("formed a ring around it"), not *They rang the city with artillery*, and *He grandstanded to the crowd* ("played to the grandstand"), not *He grandstood to the crowd.*

This principle works every time. Remember Sally Ride, the astronaut? She received a lot of publicity because she was America's first woman in space. But recently Mae Jemison did her one better. Not only is Jemison America's first *black* woman in space, but she appeared in *People* magazine in 1993 in their list of the fifty most beautiful people in the world. Publicity-wise, she has out-Sally-Rided Sally Ride (not *has out-Sally-Ridden Sally Ride*). For many years New York State's most infamous prison was Sing Sing. But since the riot at the Attica Correctional Facility in 1971, Attica has become even more infamous: it has out-Sing-Singed Sing Sing (not *has out-Sing-Sung Sing Sing*).

As for the Maple Leafs, the noun being pluralized is not *leaf*, the unit of foliage, but a noun based on the *name* Maple Leaf, Canada's national symbol. A name is not the same thing as a noun. (For example, whereas a noun may be preceded by an article like *the*, a name may not be: you cannot refer to someone as *the Donald*, unless you are Ivana Trump, whose first language is Czech.) Therefore, the noun *a Maple Leaf* (referring to, say, the goalie) must be headless, because it is a noun based on a word that is not a noun. And a noun that does not get its nounhood from one of its components cannot get an irregular plural from that component either; hence it defaults to the regular form *Maple Leafs*. This explanation also answers a question that kept bothering David Letterman thoughout one of his recent *Late Night* shows: why is the new major league baseball team in Miami called the Florida Marlins rather than the Florida Marlin, given that those fish are referred to in the plural as *marlin*? Indeed, the explanation applies to all nouns based on names:

> I'm sick of dealing with all the *Mickey Mouses* in this administration. [not *Mickey Mice*]
>
> Hollywood has been relying on movies based on comic book heroes and their sequels, like the three *Supermans* and the two *Batmans*. [not *Supermen* and *Batmen*]
>
> Why has the second half of the twentieth century produced no *Thomas Manns*? [not *Thomas Menn*]
>
> We're having Julia Child and her husband over for dinner tonight. You know, *the Childs* are great cooks. [not *the Children*]

Irregular forms, then, live at the bottom of word structure trees, where roots and stems from the mental dictionary are inserted. The

developmental psycholinguist Peter Gordon has capitalized on this effect in an ingenious experiment that shows how children's minds seem to be designed with the logic of word structure built in.

Gordon focused on a seeming oddity first noticed by the linguist Paul Kiparsky: compounds can be formed out of irregular plurals but not out of regular plurals. For example, a house infested with mice can be described as *mice-infested,* but it sounds awkward to describe a house infested with rats as *rats-infested.* We say that it is *rat-infested,* even though by definition one rat does not make an infestation. Similarly, there has been much talk about *men-bashing* but no talk about *gays-bashing* (only *gay-bashing*), and there are *teethmarks,* but no *clawsmarks.* Once there was a song about a *purple-people-eater,* but it would be ungrammatical to sing about a *purple-babies-eater.* Since the licit irregular plurals and the illicit regular plurals have similar meanings, it must be the grammar of irregularity that makes the difference.

The theory of word structure explains the effect easily. Irregular plurals, because they are quirky, have to be stored in the mental dictionary as roots or stems; they cannot be generated by a rule. Because of this storage, they can be fed into the compounding rule that joins an existing stem to another existing stem to yield a new stem. But regular plurals are not stems stored in the mental dictionary; they are complex words that are assembled on the fly by inflectional rules whenever they are needed. They are put together too late in the root-to-stem-to-word assembly process to be available to the compounding rule, whose inputs can only come out of the dictionary.

Gordon found that three- to five-year-old children obey this restriction fastidiously. Showing the children a puppet, he first asked them, "Here is a monster who likes to eat mud. What do you call him?" He then gave them the answer, *a mud-eater,* to get them started. Children like to play along, and the more gruesome the meal, the more eagerly they fill in the blank, often to the dismay of their onlooking parents. The crucial parts came next. A "monster who likes to eat mice," the children said, was a *mice-eater.* But a "monster who likes to eat rats" was never called a *rats-eater,* only a *rat-eater.* (Even the children who made the error *mouses* in their spontaneous speech never called the puppet a *mouses-eater.*) The children, in other words, respected the subtle restrictions on combining plurals and compounds inherent in the word structure rules. This suggests that the rules take the same form in the unconscious mind of the child as they do in the unconscious mind of the adult.

But the most interesting discovery came when Gordon examined how children might have acquired this constraint. Perhaps, he reasoned, they learned it from their parents by listening for whether the plurals that occur inside the parents' compounds are irregular, regular, or both, and then duplicate whatever kinds of compounds they hear. This would be impossible, he discovered. Motherese just doesn't have any compounds containing plurals. Most compounds are like *toothbrush,* with singular nouns inside them; compounds like *mice-infested,* though grammatically possible, are seldom used. The children produced *mice-eater* but never *rats-eater,* even though they had no evidence from adult speech that this is how languages work. We have another demonstration of knowledge despite "poverty of the input," and it suggests that another basic aspect of grammar may be innate. Just as Crain and Nakayama's Jabba experiment showed that in syntax children automatically distinguish between word strings and phrase structures, Gordon's mice-eater experiment shows that in morphology children automatically distinguish between roots stored in the mental dictionary and inflected words created by a rule.

A word, in a word, is complicated. But then what in the world *is* a word? We have just seen that "words" can be built out of parts by morphological rules. But then what makes them different from phrases or sentences? Shouldn't we reserve the word "word" for a thing that has to be rote-memorized, the arbitrary Saussurean sign that exemplifies the first of the two principles of how language works (the other being the discrete combinatorial system)? The puzzlement comes from the fact that the everyday word "word" is not scientifically precise. It can refer to two things.

The concept of a word that I have used so far in this chapter is a linguistic object that, even if built out of parts by the rules of morphol-

ogy, behaves as the indivisible, smallest unit with respect to the rules of syntax—a "syntactic atom," in *atom*'s original sense of something that cannot be split. The rules of syntax can look inside a sentence or phrase and cut and paste the smaller phrases inside it. For example, the rule for producing questions can look inside the sentence *This monster eats mice* and move the phrase corresponding to *mice* to the front, yielding *What did this monster eat?* But the rules of syntax halt at the boundary between a phrase and a word; even if the word is built out of parts, the rules cannot look "inside" the word and fiddle with those parts. For example, the question rule cannot look inside the word *mice-eater* in the sentence *This monster is a mice-eater* and move the morpheme corresponding to *mice* to the front; the resulting question is virtually unintelligible: *What is this monster an -eater?* (Answer: mice.) Similarly, the rules of syntax can stick an adverb inside a phrase, as in *This monster eats mice quickly.* But they cannot stick an adverb inside a word, as in *This monster is a mice-quickly-eater.* For these reasons, we say that words, even if they are generated out of parts by one set of rules, are not the same thing as phrases, which are generated out of parts by a different set of rules. Thus one precise sense of our everyday term "word" refers to the units of language that are the products of morphological rules, and which are unsplittable by syntactic rules.

The second, very different sense of "word" refers to a rote-memorized chunk: a string of linguistic stuff that is arbitrarily associated with a particular meaning, one item from the long list we call the mental dictionary. The grammarians Anna Maria Di Sciullo and Edwin Williams coined the term "listeme," the unit of a memorized list, to refer to this sense of "word" (their term is a play on "morpheme," the unit of morphology, and "phoneme," the unit of sound). Note that a listeme need not coincide with the first precise sense of "word," a syntactic atom. A listeme can be a tree branch any size, as long as it cannot be produced mechanically by rules and therefore has to be memorized. Take idioms. There is no way to predict the meaning of *kick the bucket, buy the farm, spill the beans, bite the bullet, screw the pooch, give up the ghost, hit the fan,* or *go bananas* from

the meanings of their components using the usual rules of heads and role-players. *Kicking the bucket* is not a kind of kicking, and buckets have nothing to do with it. The meanings of these phrase-sized units have to be memorized as listemes, just as if they were simple word-sized units, and so they are really "words" in this second sense. Di Sciullo and Williams, speaking as grammatical chauvinists, describe the mental dictionary (lexicon) as follows: "If conceived of as the set of listemes, the lexicon is incredibly boring by its very nature. . . . The lexicon is like a prison—it contains only the lawless, and the only thing that its inmates have in common is their lawlessness."

In the rest of this chapter I turn to the second sense of "word," the listeme. It will be a kind of prison reform: I want to show that the lexicon, though a repository of lawless listemes, is deserving of respect and appreciation. What seems to a grammarian like an act of brute force incarceration—a child hears a parent use a word and thenceforth retains that word in memory—is actually an inspiring feat.

One extraordinary feature of the lexicon is the sheer capacity for memorization that goes into building it. How many words do you think an average person knows? If you are like most writers who have offered an opinion based on the number of words they hear or read, you might guess a few hundred for the uneducated, a few thousand for the literate, and as many as 15,000 for gifted wordsmiths like Shakespeare (that is how many distinct words are found in his collected plays and sonnets).

The real answer is very different. People can recognize vastly more words than they have occasion to use in some fixed period of time or space. To estimate the size of a person's vocabulary—in the sense of memorized listemes, not morphological products, of course, because the latter are infinite—psychologists use the following method. Start with the largest unabridged dictionary available; the smaller the dictionary, the more words a person might know but not get credit for. Funk & Wagnall's *New Standard Unabridged Dictionary,* to take an example, has 450,000 entries, a healthy number, but too many to test exhaustively. (At thirty seconds a word, eight hours

a day, it would take more than a year to test a single person.) Instead, draw a sample—say, the third entry from the top of the first column on every eighth left-hand page. Entries often have many meanings, such as "*hard*: (1) firm; (2) difficult; (3) harsh; (4) toilsome . . ." and so on, but counting them would require making arbitrary decisions about how to lump or split the meanings. Thus it is practical only to estimate how many words a person has learned at least one meaning for, not how many meanings a person has learned altogether. The testee is presented with each word in the sample, and asked to choose the closest synonym from a set of alternatives. After a correction for guessing, the proportion correct is multiplied by the size of the dictionary, and that is an estimate of the person's vocabulary size.

Actually, another correction must be applied first. Dictionaries are consumer products, not scientific instruments, and for advertising purposes their editors often inflate the number of entries. ("Authoritative. Comprehensive. Over 1.7 million words of text and 160,000 definitions. Includes a 16-page full-color atlas.") They do it by including compounds and affixed forms whose meanings are predictable from the meanings of their roots and the rules of morphology, and thus are not true listemes. For example, my desk dictionary includes, together with *sail*, the derivatives *sailplane, sailer, sailless, sailing-boat,* and *sailcloth,* whose meanings I could deduce even if I had never heard them before.

The most sophisticated estimate comes from the psychologists William Nagy and Richard Anderson. They began with a list of 227,553 different words. Of these, 45,453 were simple roots and stems. Of the remaining 182,100 derivatives and compounds, they estimated that all but 42,080 could be understood in context by someone who knew their components. Thus there were a total of 44,453 + 42,080 = 88,533 listeme words. By sampling from this list and testing the sample, Nagy and Anderson estimated that an average American high school graduate knows 45,000 words—three times as many as Shakespeare managed to use! Actually, this is an underestimate, because proper names, numbers, foreign words, acronyms, and many common undecomposable compounds were excluded. There is

no need to follow the rules of Scrabble in estimating vocabulary size; these forms are all listemes, and a person should be given credit for them. If they had been included, the average high school graduate would probably be credited with something like 60,000 words (a tetrabard?), and superior students, because they read more, would probably merit a figure twice as high, an octobard.

Is 60,000 words a lot or a little? It helps to think of how quickly they must have been learned. Word learning generally begins around the age of twelve months. Therefore, high school graduates, who have been at it for about seventeen years, must have been learning an average of ten new words a day continuously since their first birthdays, or about a new word every ninety waking minutes. Using similar techniques, we can estimate that an average six-year-old commands about 13,000 words (notwithstanding those dull, dull *Dick and Jane* reading primers, which were based on ridiculously lowball estimates). A bit of arithmetic shows that preliterate children, who are limited to ambient speech, must be lexical vacuum cleaners, inhaling a new word every two waking hours, day in, day out. Remember that we are talking about listemes, each involving an arbitrary pairing. Think about having to memorize a new batting average or treaty date or phone number every ninety minutes of your waking life since you took your first steps. The brain seems to be reserving an especially capacious storage space and an especially rapid transcribing mechanism for the mental dictionary. Indeed, naturalistic studies by the psychologist Susan Carey have shown that if you casually slip a new color word like *olive* into a conversation with a three-year-old, the child will probably remember something about it five weeks later.

Now think of what goes into each act of memorization. A word is the quintessential symbol. Its power comes from the fact that every member of a linguistic community uses it interchangeably in speaking and understanding. If you use a word, then as long as it is not too obscure I can take it for granted that if I later utter it to a third party, he will understand my use of it the same way I understood yours. I do not have to try the word back on you to see how you react, or test it out

on every third party and see how they react, or wait for you to use it with third parties. This sounds more obvious than it is. After all, if I observe that a bear snarls before it attacks, I cannot expect to scare a mosquito by snarling at it; if I bang a pot and the bear flees, I cannot expect the bear to bang a pot to scare hunters. Even within our species, learning a word from another person is not just a case of imitating that person's behavior. Actions are tied to particular kinds of actors and targets of the action in ways that words are not. If a girl learns to flirt by watching her older sister, she does not flirt with the sister or with their parents but only with the kind of person that she observes to be directly affected by the sister's behavior. Words, in contrast, are a universal currency within a community. In order to learn to use a word upon merely hearing it used by others, babies must tacitly assume that a word is not merely a person's characteristic behavior in affecting the behavior of others, but a shared bidirectional symbol, available to convert meaning to sound by any person when the person speaks, and sound to meaning by any person when the person listens, according to the same code.

Since a word is a pure symbol, the relation between its sound and its meaning is utterly arbitrary. As Shakespeare (using a mere tenth of a percent of his written lexicon and a far tinier fraction of his mental one) put it,

> What's in a name? that which we call a rose
> By any other name would smell as sweet.

Because of that arbitrariness, there is no hope that mnemonic tricks might lighten the memorization burden, at least for words that are not built out of other words. Babies should not, and apparently do not, expect *cattle* to mean something similar to *battle*, or *singing* to be like *stinging*, or *coats* to resemble *goats*. Onomatopoeia, where it is found, is of no help, because it is almost as conventional as any other word sound. In English, pigs go "oink"; in Japanese, they go "boo-boo." Even in sign languages the mimetic abilities of the hands are put aside and their configurations are treated as arbitrary symbols. Residues of resemblance between a sign and its referent can occasionally be discerned, but like onomatopoeia they are so much in the eye or ear of the beholder that they are of little use in learning. In American Sign Language the sign for "tree" is a motion of a hand as if it was a branch waving in the wind; in Chinese Sign Language "tree" is indicated by the motion of sketching a tree trunk.

The psychologist Laura Ann Petitto has a startling demonstration that the arbitrariness of the relation between a symbol and its meaning is deeply entrenched in the child's mind. Shortly before they turn two, English-speaking children learn the pronouns *you* and *me*. Often they reverse them, using *you* to refer to themselves. The error is forgivable. *You* and *me* are "deictic" pronouns, whose referent shifts with the speaker: *you* refers to you when I use it but to me when you use it. So children may need some time to get that down. After all, Jessica hears her mother refer to her, Jessica, using *you;* why should she not think that *you* means "Jessica"?

Now, in ASL the sign for "me" is a point to one's chest; the sign for "you" is a point to one's partner. What could be more transparent? One would expect that using "you" and "me" in ASL would be as foolproof as knowing how to point, which all babies, deaf and hearing, do before their first birthday. But for the deaf children Petitto studied, pointing is not pointing. The children used the sign of pointing to their conversational partners to mean "me" at exactly the age at which hearing children use the spoken sound *you* to mean "me." The children were treating the gesture as a pure linguistic symbol; the fact that it pointed somewhere did not register as being relevant. This attitude is appropriate in learning sign languages; in ASL, the pointing hand-shape is like a meaningless consonant or vowel, found as a component of many other signs, like "candy" and "ugly."

There is one more reason we should stand in awe of the simple act of learning a word. The logician W. V. O. Quine asks us to imagine a linguist studying a newly discovered tribe. A rabbit scurries by, and a native shouts, "Gavagai!" What does *gavagai* mean? Logically speaking, it needn't be "rabbit." It could refer to that particular rabbit (Flopsy, for example). It could mean any furry thing, any mammal, or

any member of that species of rabbit (say, *Oryctolagus cuniculus*), or any member of that variety of that species (say, chinchilla rabbit). It could mean scurrying rabbit, scurrying thing, rabbit plus the ground it scurries upon, or scurrying in general. It could mean footprint-maker, or habitat for rabbit-fleas. It could mean the top half of a rabbit, or rabbit-meat-on-the-hoof, or possessor of at least one rabbit's foot. It could mean anything that is either a rabbit or a Buick. It could mean collection of undetached rabbit parts, or "Lo! Rabbithood again!," or "It rabbiteth," analogous to "It raineth."

The problem is the same when the child is the linguist and the parents are the natives. Somehow a baby must intuit the correct meaning of a word and avoid the mind-boggling number of logically impeccable alternatives. It is an example of a more general problem that Quine calls "the scandal of induction," which applies to scientists and children alike: how can they be so successful at observing a finite set of events and making some correct generalization about all future events of that sort, rejecting an infinite number of false generalizations that are also consistent with the original observations?

We all get away with induction because we are not open-minded logicians but happily blinkered humans, innately constrained to make only certain kinds of guesses—the probably correct kinds—about how the world and its occupants work. Let's say the word-learning baby has a brain that carves the world into discrete, bounded, cohesive objects and into the actions they undergo, and that the baby forms mental categories that lump together objects that are of the same kind. Let's also say that babies are designed to expect a language to contain words for kinds of objects and words for kinds of actions—nouns and verbs, more or less. Then the undetached rabbit parts, rabbit-trod ground, intermittent rabbiting, and other accurate descriptions of the scene will, fortunately, not occur to them as possible meanings of *gavagai*.

But could there really be a preordained harmony between the child's mind and the parent's? Many thinkers, from the woolliest mystics to the sharpest logicians, united only in their assault on common sense, have claimed that the distinction between an object and an action is not in the world or even in our minds, initially, but is imposed on us by our language's distinction between nouns and verbs. And if it is the word that delineates the thing and the act, it cannot be the concepts of thing and act that allow for the learning of the word.

I think common sense wins this one. In an important sense, there really are things and kinds of things and actions out there in the world, and our mind is designed to find them and to label them with words. That important sense is Darwin's. It's a jungle out there, and the organism designed to make successful predictions about what is going to happen next will leave behind more babies designed just like it. Slicing space-time into objects and actions is an eminently sensible way to make predictions given the way the world is put together. Conceiving of an extent of solid matter as a thing—that is, giving a single mentalese name to all of its parts—invites the prediction that those parts will continue to occupy some region of space and will move as a unit. And for many portions of the world, that prediction is correct. Look away, and the rabbit still exists; lift the rabbit by the scruff of the neck, and the rabbit's foot and the rabbit ears come along for the ride.

What about kinds of things, or categories? Isn't it true that no two individuals are exactly alike? Yes, but they are not arbitary collections of properties, either. Things that have long furry ears and tails like pom-poms also tend to eat carrots, scurry into burrows, and breed like, well, rabbits. Lumping objects into categories—giving them a category label in mentalese—allows one, when viewing an entity, to infer some of the properties one cannot directly observe, using the properties one *can* observe. If Flopsy has long furry ears, he is a "rabbit"; if he is a rabbit, he might scurry into a burrow and quickly make more rabbits.

Moreover, it pays to give objects several labels in mentalese, designating different-sized categories like "cottontail rabbit," "rabbit," "mammal," "animal," and "living thing." There is a tradeoff involved in choosing one category over another. It takes less effort to determine that Peter Cottontail is an animal than that he is a cotton-

tail (for example, an animallike motion will suffice for us to recognize that he is an animal, leaving it open whether or not he is a cottontail). But we can predict more new things about Peter if we know he is a cottontail than if we merely know he is an animal. If he is a cottontail, he likes carrots and inhabits open country or woodland clearings; if he is merely an animal, he could eat anything and live anywhere, for all one knows. The middle-sized or "basic-level" category "rabbit" represents a compromise between how easy it is to label something and how much good the label does you.

Finally, why separate the rabbit from the scurry? Presumably because there are predictable consequences of rabbithood that cut across whether it is scurrying, eating, or sleeping: make a loud sound, and in all cases it will be down a hole lickety-split. The consequences of making a loud noise in the presence of lionhood, whether eating or sleeping, are predictably different, and that is a difference that makes a difference. Likewise, scurrying has certain consequences regardless of who is doing it; whether it be rabbit or lion, a scurrier does not remain in the same place for long. With sleeping, a silent approach will generally work to keep a sleeper—rabbit or lion—motionless. Therefore a powerful prognosticator should have separate sets of mental labels for kinds of objects and kinds of actions. That way, it does not have to learn separately what happens when a rabbit scurries, what happens when a lion scurries, what happens when a rabbit sleeps, what happens when a lion sleeps, what happens when a gazelle scurries, what happens when a gazelle sleeps, and on and on; knowing about rabbits and lions and gazelles in general, and scurrying and sleeping in general, will suffice. With m objects and n actions, a knower needn't go through $m \times n$ learning experiences; it can get away with $m + n$ of them.

So even a wordless thinker does well to chop continuously flowing experience into things, kinds of things, and actions (not to mention places, paths, events, states, kinds of stuff, properties, and other types of concepts). Indeed, experimental studies of baby cognition have shown that infants have the concept of an object before they learn any words for objects, just as we would expect. Well before their

first birthday, when first words appear, babies seem to keep track of the bits of stuff that we would call objects: they show surprise if the parts of an object suddenly go their own ways, or if the object magically appears or disappears, passes through another solid object, or hovers in the air without visible means of support.

Attaching words to these concepts, of course, allows one to share one's hard-won discoveries and insights about the world with the less experienced or the less observant. Figuring out which word to attach to which concept is the *gavagai* problem, and if infants start out with concepts corresponding to the kinds of meanings that languages use, the problem is partly solved. Laboratory studies confirm that young children assume that certain kinds of concepts get certain types of words, and other kinds of concepts cannot be the meaning of a word at all. The developmental psychologists Ellen Markman and Jeanne Hutchinson gave two- and three-year-old children a set of pictures, and for each picture asked them to "find another one that is the same as this." Children are intrigued by objects that interact, and when faced with these instructions they tend to select pictures that make groups of role-players like a blue jay and a nest or a dog and a bone. But when Markman and Hutchinson told them to "find another *dax* that is the same as this *dax*," the children's criterion shifted. A word must label a *kind* of thing, they seemed to be reasoning, so they put together a bird with another type of bird, a dog with another type of dog. For a child, a *dax* simply cannot mean "a dog or its bone," interesting though the combination may be.

Of course, more than one word can be applied to a thing: Peter Cottontail is not only a *rabbit* but an *animal* and a *cottontail*. Children have a bias to interpret nouns as middle-level kinds of objects like "rabbit," but they also must overcome that bias, to learn other types of words like *animal*. Children seem to manage this by being in sync with a striking feature of language. Though most common words have many meanings, few meanings have more than one word. That is, homonyms are plentiful, synonyms rare. (Virtually all supposed synonyms have some difference in meaning, however small. For example, *skinny* and *slim* differ in their connotation of desirability; *police-*

man and *cop* differ in formality.) No one really knows why languages are so stingy with words and profligate with meanings, but children seem to expect it (or perhaps it is this expectation that causes it!), and that helps them further with the *gavagai* problem. If a child already knows a word for a kind of thing, then when another word is used for it, he or she does not take the easy but wrong way and treat it as a synonym. Instead, the child tries out some other possible concept. For example, Markman found that if you show a child a pair of pewter tongs and call it *biff,* the child interprets *biff* as meaning tongs in general, showing the usual bias for middle-level objects, so when asked for "more biffs," the child picks out a pair of plastic tongs. But if you show the child a pewter cup and call it *biff,* the child does not interpret *biff* as meaning "cup," because most children already know a word that means "cup," namely, *cup*. Loathing synonyms, the children guess that *biff* must mean something else, and the stuff the cup is made of is the next most readily available concept. When asked for more *biffs,* the child chooses a pewter spoon or pewter tongs.

Many other ingenious studies have shown how children home in on the correct meanings for different kinds of words. Once children know some syntax, they can use it to sort out different kinds of meaning. For example, the psychologist Roger Brown showed children a picture of hands kneading a mass of little squares in a bowl. If he asked them, "Can you see any sibbing?," the children pointed to the hands. If instead he asked them, "Can you see a sib?," they point to the bowl. And if he asked, "Can you see any sib?," they point to the stuff inside the bowl. Other experiments have uncovered great sophistication in children's understanding of how classes of words fit into sentence structures and how they relate to concepts and kinds.

So what's in a name? The answer, we have seen, is, a great deal. In the sense of a morphological product, a name is an intricate structure, elegantly assembled by layers of rules and lawful even at its quirkiest. And in the sense of a listeme, a name is a pure symbol, part of a cast of thousands, rapidly acquired because of a harmony between the mind of the child, the mind of the adult, and the texture of reality.

6

The Sounds of Silence

When I was a student I worked in a laboratory at McGill University that studied auditory perception. Using a computer, I would synthesize trains of overlapping tones and determine whether they sounded like one rich sound or two pure ones. One Monday morning I had an odd experience: the tones suddenly turned into a chorus of screaming munchkins. Like this: (beep boop-boop) (beep boop-boop) (beep boop-boop) HUMPTY-DUMPTY-HUMPTY-DUMPTY-HUMPTY-DUMPTY (beep boop-boop) (beep boop-boop) HUMPTY-DUMPTY-HUMPTY-DUMPTY-HUMPTY-HUMPTY-DUMPTY-DUMPTY (beep boop-boop) (beep boop-boop) (beep boop-boop) HUMPTY-DUMPTY (beep boop-boop) HUMPTY-HUMPTY-HUMPTY-DUMPTY (beep boop-boop). I checked the oscilloscope: two streams of tones, as programmed. The effect had to be perceptual. With a bit of effort I could go back and forth, hearing the sound as either beeps or munchkins. When a fellow student entered, I recounted my discovery, mentioning that I couldn't wait to tell Professor Bregman, who directed the laboratory. She offered some advice: don't tell anyone, except perhaps Professor Poser (who directed the psychopathology program).

Years later I discovered what I had discovered. The psychologists