Chapter 7
New to a child, new to English: Learning and introducing words /ˈnjuw tuw ə ˈtʃajld, ˈnjuw tuw ˈɪŋglɪʃ: ˈlʌɹnɪŋ ən ˌɪntrəˈduwsɪŋ ˈwəɹdz /

In this chapter, we look first at what kinds of assumptions children have to be making in order to figure out their meaning from their observed uses. Children have to be able to recognize objects, make guesses about what the people around them are talking about, and assume that concepts have only one name. It turns out that some of the most valuable clues about content words' meanings come from the function words they combine with. We then move on to consider where new words come from, examining processes of clipping, acronym-formation, blends, and compounding.

7.1 How do children learn the meanings of words?

When you're listening to someone speak a language that you don't know, you can't even tell where the individual phonological words begin and end. As we saw at the end of chapter 3, babies are in this state initially, but over time, as they hear more and more speech, they come to recognize the phonotactic patterns of English and begin to parse the speech stream into phonological words. They begin to track and remember recurring patterns. Before long, they have an inventory of highly frequent

phonological sequences in their memory—morphemes. That's a pretty amazing accomplishment in itself. But now they have an even steeper problem to solve.

As we saw in chapter 1, the sound-meaning relationship is arbitrary. That is, the particular phonemes a morpheme is made out of don't provide any hint about what the morpheme might mean. The child has to figure out the meaning of individual morphemes and groups of morphemes on his own—he has to incorporate his statistically-discovered morphemes into listemes. How can he do this? How do children learn the meanings of words?

Children, and people generally, are phenomenally good at it. The average young adult has command of between 50,000 and 100,000 listemes, counting affixes, roots, and word-sized and phrase-sized idioms. Babies start producing words on their at around one year of age. On average, they start off fairly slowly, learning about one word every three days. At about 18 months, they're producing a new word almost every day. Then comes the real spurt: between the ages of 2 and 6, they're up to four new words a day, and by the time they're ten years old, they're learning about twelve new words a day. The reason that most people top out between 50,000 and 100,000 listemes is because they stop running into new ones on a regular basis, not because they lose their ability to learn new words. The more words you encounter, the more you will learn.

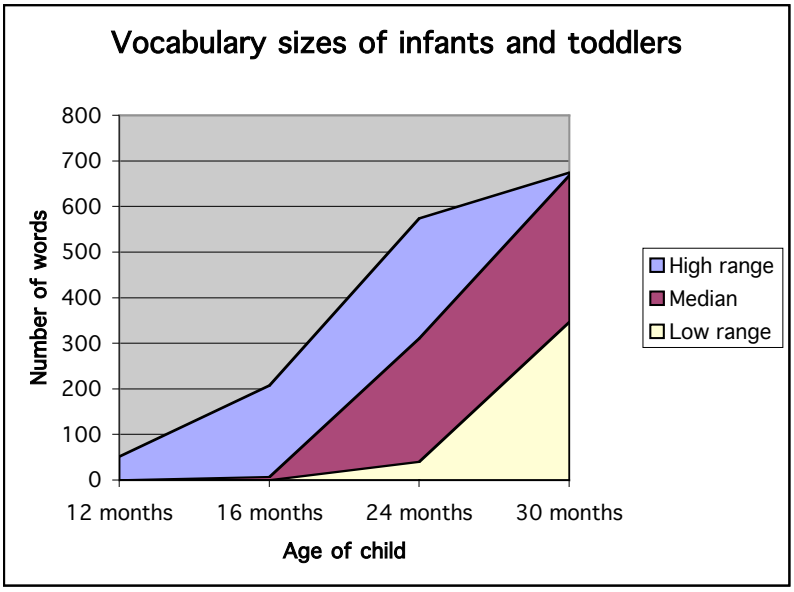

So the average 6-month old baby has command of no listemes, and the average 18 year old can deploy about 50,000. In between, she learned it all, mostly without noticing. Only very occasionally is a learner explicitly instructed about listeme meaning (usually when she uses the listeme in a way that reveals that she has made the wrong hypothesis about what the meaning is.) For the most part, the learner comes up with a meaning for a novel word by simple induction: given a linguistic sound sequence and a context, the learner assigns a meaning to the sound sequence.

Let's be sure we understand why this is a hard problem. For a long time, people didn't really think it *was* a hard problem at all. As soon as anyone wondered about how children came to attach certain sound combinations to certain meanings, the association hypothesis was born.

The philosopher Locke (1690) assumed that all babies had to do to learn a word was notice co-occurrences of sound sequences and their referents. If a baby saw a rabbit go by, and Mother said /ˈɹæbɪt/ at that moment, and that happened more than once, eventually the baby would come to associate the sound /ˈɹæbɪt/ with real-world rabbits, and would hence know the meaning of the listeme. This general picture fit in well

with behaviorist models of learning. After all, if you ring a bell before feeding them enough times, dogs will involuntarily begin to salivate at the sound of bells, as Pavlov found out. Why shouldn't word learning work the same way? Pronounce /ˈɹæbɨt/ in the presence of an actual rabbit enough times, and a child will involuntarily begin thinking of a rabbit when he hears the sound /ˈɹæbɨt/. Better yet, if some reward is given for demonstrating the correct sound-meaning association, like a smile from Mom, or getting a bottle when /ˈbɑbɑ/ is pronounced, the positive reinforcement will help the associationist mechanism operate even faster.

It seems clear that sound-object association must play a role in children's word learning to a certain extent. It's equally clear, however, that it can't possibly be the whole story.

To consider the most basic kind of objection first, let's look at the following quote from Quine discussing the problem of matching up sound sequences and real-world items:

> "The recovery of a man's current language from his currently observed responses is the task of the linguist who, unaided by an interpreter, is out to penetrate and translate a language hitherto unknown. All the objective data he has to go on are the forces that he sees impinging on the native's surfaces and the observable behavior, vocal and otherwise, of the native.... A rabbit scurries by, the native says 'Gavagai', and the linguist notes down the sentence 'Rabbit' (or 'Lo, a rabbit')... Who knows but what the objects to which this term applies are not rabbits at all, but mere stages, or brief temporal segments, of rabbits? In either event the stimulus situations that prompt assent to 'Gavagai' would be the same as for 'Rabbit'. Or perhaps the objects to which 'gavagai' applies are all and sundry undetached parts of rabbits; again the stimulus meaning would register no difference. When from the sameness of stimulus meaning, the linguist leaps to the conclusion that a gavagai is a whole enduring rabbit, he is just taking for granted that the native is enough like us to have a brief general term for rabbits and no brief general term for rabbit stages or parts. A further alternative...is to take 'gavagai' as a singular term naming the fusion, in Goodman's sense, of all rabbits; that single though discontinuous partion of the spatiotemporal world that consists of rabbits.... And a still further alternative in the case of 'gavagai' is to take it as a singular term naming a recurring universal: rabbithood. The distinction between concrete and abstract object, as well as that between general and singular term, is independent of stimulus meaning.'
>
> *Quine 1960: 29, 52*

Substitute 'child' for 'linguist' and 'mother' for 'native' in the discussion above, and the discussion is exactly applicable to first-language acquisition. Quine's point is that there is an infinite number of things in any given environment that a sound sequence could be associated with. Even tracking the occurrence of a sound sequence across several different occurrences in several different environments will never eliminate all the alternative available referents. The string 'gavagai' could refer to anything—even, for instance, something about the person listening to the word—and it's in principle impossible for a hearer to proceed by remembering all the details of all the circumstances in which a word was heard, cross-classify the circumstances, and come up with the one element common to them all that the word must refer to. Word learners—babies—must be making some additional assumptions that let them narrow their hypothesis space a bit.

7.2 Learning words for middle-sized observables

Obviously, in any given encounter with a new sound sequence, a baby *doesn't* consider *every* possible association with that sound sequence. They make educated guesses about what's the most likely association, given what else they know about the situation. Then the question becomes, what else do they know? What assumptions are they basing their educated guesses on?

Experimenters have gotten very good at inferring what assumptions babies are making about the way the world works. Many of the experiments that have been conducted to investigate this problem are based on the length of time a baby looks at a scene, or on the rate of pacifier-sucking, like the ones described in chapter 3. The basic idea is that longer looking at a scene, or increased rates of pacifier-sucking, indicate feelings of surprise at or interest in something in the scene. Surprise and interest are the result when things don't happen the way the baby expects them to—so they reveal what the baby expects, i.e. what its basic assumptions are.

7.2.1 Whole-object bias

All else being equal, babies tend to assume that novel sound sequences refer to entire objects (like rabbits), not to some subpart of the object (like ears), or to some property of the object (like softness), but to the whole object itself. Experimenters have determined that babies automatically assume several things about objects: they're not attached to their surroundings, they move as a solid unit rather than piece by piece, they must be supported from below, they have a distinctive shape, etc. If you construct a situation where an apparently object-like item violates any of these basic assumptions, even very young babies will exhibit surprise. They know how objects work. Furthermore, they prefer to assume that novel sound sequences name objects. They'll move away from that assumption if other clues contradict it (see below), but for an initial guess, Locke's hypothetical baby would be likely to prefer the rabbit itself as the referent of /ˈræbɨt/, rather than its parts, its properties, or some element of the background. The whole-object bias is what Quine is referring to when he says, "[the linguist] is just taking for granted that the native is enough like us to have a brief general term for rabbits and no brief general term for rabbit stages or parts." That's what babies take for granted too.

7.2.2 Mind-reading bias

The best clues about what a particular utterance is intended to refer to come not from the environment, but from the utterer herself. If the person speaking is referring to an object in the immediate domain, her gaze will tend to go to that object during the utterance. If a baby can learn to follow eye gaze, and infer that the person he's listening to is thinking about the thing that her eyes are gazing at, he's got a great clue about what the utterance is intended to refer to. If the mother of Locke's hypothetical baby were looking right at the rabbit when she said /ˈræbɨt/, and if the baby noticed where her eyes were looking at the time, he'd have one more reason to prefer the actual rabbit as the referent of /ˈræbɨt/, rather than any other object in the environment.

Learning to follow eye gaze and make inferences about the gazer's thoughts is a very advanced cognitive skill. Not many animals can do it, if

any. Babies learn to do it between about 12 and 14 months of age, right around the time that they begin producing words for the first time.

7.2.3 Mutual exclusivity bias

Imagine that there's a mouse very close by the rabbit. The child can't tell which of the two the speaker is looking at, and they're both perfectly good and equally interesting objects. How can he choose between these two possible referents for /ˈræbɨt/?

Well, imagine that he already knows the word for the mouse—he's already, on the basis of former experience with mice when no rabbits were around, associated /ˈmaws/ with mice. But he's never seen a rabbit before. It turns out that in situations like this, children assume the new sound corresponds to the novel object. That is, unless there is good evidence to the contrary, they won't imagine that a single object can correspond to two different sounds sequences. Because they know the mouse goes with /ˈmaws/, they'll assume that /ˈræbɨt/ can't go with the mouse. Consequently, it must go with the rabbit. This is called the mutual exclusivity bias—children assume that different sound sequences have to go with different referents.

7.2.4 The taxonomic and meronymic biases

The mutual exclusivity bias gives kids a leg up in trying to figure out the cases where a different sound sequence *is* applied to an object they already know a word for. Let's say that Locke's hypothetical baby, who knows the word *mouse* already, is hanging around with his mother when a mouse comes by, and she looks at it and says /ˈɹowdənt/, or /ˈgrej/, or /ˈijɹ/. He *knows* she has to be referring to the mouse, but she's not using the word he's familiar with. Mutual exclusivity rules out the possibility that the new sound sequence might mean 'mouse'. What else might it mean?

It might be a word that names a type of object that mice are an instance of, but not the only instance of—it might name a bigger class of objects than 'mouse' names. This would be the right guess in the case of /ˈɹowdənt/—a mouse is a kind of rodent, although there are others. If mutual exclusivity tells the baby to abandon the whole-object hypothesis,

they could adopt this *taxonomic* hypothesis, and get the right answer sometimes.

Alternatively, it might be a word that names a subpart of the mouse—something that comes along with mice, but which doesn't make up the whole mouse by itself. This would be the right guess in the case of /ˈijɹ/. If mutual exclusivity tells the baby to abandon the whole-object hypothesis, they could adopt this *meronymic* hypothesis.

Some interesting experiments have shown that the taxonomic bias is specific to the word-learning situation. Shown a picture of a cow, a pig, and a glass of milk, kids (who know the words for all three things) will tend to put the cow and the milk together if just asked to sort the objects into groups—the cow and the milk are naturally associated with each other, and the association forms the basis for the sorting. If, on the other hand, the experimenter points to the cow, says something like *That's a dax*, and then asks the kids to sort the pictures into groups of daxes and non-daxes, the kids will tend to put the cow and the pig together, rather than the cow and the milk. That is, they have assumed that *dax* must mean something like 'animal', rather than something like 'cow-related items'—even though their previous sorting showed that they are predisposed to sort according to cow-related items and non-cow-related items. Kids apparently assume that word meanings are related in taxonomic ways, rather than in purely associative ways.

7.2.5 When the basics fail

What if the mother of Locke's hypothetical baby, who was faced with a rabbit for the first time ever, *hadn't* said 'rabbit'? What if she had been thinking about animals in general, and had said /ˈænɨməl/? What if she'd been thinking about ears, or fur, and said /ˈijr/ or /ˈfʌɹ/? What if she knew the rabbit was named Peter, and had said /ˈpijtəɹ/? Her eye-gaze clue would have been the same, and the whole-object bias would operate in the same way. The exclusivity bias wouldn't apply, because the baby had never seen the rabbit before, or learned the meaning of the sound sequence. Given everything we've said, the baby would have to assume that the new sound sequence referred to rabbits—and he'd be wrong.

In fact, kids do occasionally make this kind of mistake. There are many stories of children who assumed that *rover* was the word for all dogs, not just the name of the family pet—or vice versa: kids who've assumed that the family pet's name is Dog. But they don't make it as often as you might think. The kids are exploiting other clues to inform their guesses.

7.2.6 Morphological and syntactic clues

There's a crucial ingredient missing from this puzzle, which is that in fact Locke's hypothetical baby's mother is very unlikely to just say 'rabbit' or 'ear' or 'fur' or even 'Peter' as an isolated utterance. Nearly every content word a baby hears will be embedded in a sentence, like *There goes the rabbit!* or *Do you see Peter*? or *Look how long his ears are!* It's almost never the case that an utterance will consist of just a single content word.

Remember that kids are not just trying to figure out the meanings of content words this whole time. Their pattern-matching statistics will have zeroed in on the *function* morphemes as well—determiners, plural marking, pronouns, complementizers, conjunctions—the lot. In fact, because the function morphemes recur so much more frequently than content words, kids are very likely to have isolated them as units very early on. In the Brown corpus of English—a million words of English text—the first 62 most frequent words are function words. The most frequent content word that appears in the corpus is ranked 63rd in frequency. (It's *new*). The top ten most frequent words in the corpus, in order, are *the, of, and, to, a, in, that, is, was, he*. Of course, the frequency count on a corpus like this doesn't take into account non-word listemes, but we can be sure that affixes like *-s, -ed* and *-ing* would appear in the top ranks if they were included. Certainly children's statistical analyzers will detect and remember not only stand-alone function listemes like *to* and *the*, but also affixal function listemes. Having isolated the function words early, language-learning children are bound to start noticing when they show up and when they don't.

Even if they haven't attached any meaning to the function words yet, for instance, they may have noticed that different content words co-occur with different function words. Here are two very similar utterances

that our hypothetical mother might have made in the presence of the rabbit:

(115) a. It′s a rabbit!
b. It′s Peter!

Our hypothetical baby will notice whether the function morpheme *a* is present or absent. His statistical tallies will have told him that *a* can co-occur with some content words but not others—*girl, bottle, apple*, but not *Jimmy, Susan,* or *juice*. He can immediately put /ˈɹæbɨt/ in the same class as the other words he's heard with *a*, and if he's learned the meanings of some of those words already, he might guess that /ˈɹæbɨt/ names the same kind of concept that those other words do. Since *a* goes with count nouns, and object names are the quintessential examples of count nouns, he might find that the object-name guess about the meaning of /ˈɹæbɨt/ was reinforced. In contrast, since there's no *a* (or any other determiner) with the word Peter, if he heard the utterance in (115)b above he might be able to conclude that /ˈpijtəɹ/ is *not* a count noun. It would take a little more evidence to decide what kind of thing it was—it might be a mass noun (*It's fur!*) or an adjective (*It's grey!*) as well as a proper name—but at least the absence of a determiner would steer the baby away from a count-noun-like meaning for /ˈpijtəɹ/.

7.3 Learning words for non-observables

It should be clear by now that word learning isn't a trivial problem, even word-learning of words for everyday middle-sized objects. In order to winnow down the myriad possibilities for the meaning of any new sound sequence, a child must be making several quite strong tacit assumptions. That is, when learning words for middle-sized observables, children have to be applying some fairly structured cognitive principles to the problem—it can't just be straight sound-stimulus association.

What's even worse is that Mom may even say /ˈræbɨt/ when no rabbits are around! Quine discusses the problems that this situation poses for his imaginary linguist:

> The difficulty is that an informant's assent to or dissent from 'Gavagai' can depend excessively on prior collateral information... He may assent on the occasion of nothing better than an ill-glimpsed movement in the grass, because of his earlier observation, unknown to the linguist, of rabbits near the spot. ... More persistent discrepancies of the same type can be imagined, affecting not one native but all, and not once but regularly. There may be a local rabbit-fly, unknown to the linguist, and recognizable some way off by its long wings and erratic movements; and seeing such a fly in the neighborhood of an ill-glimpsed animal could help a native to recognize the latter as a rabbit.

How can a child know in advance when the referent of /ˈɹæbɨt/ is present in his field of view and when it's not? The association between utterances of a word and the physical presence of its referent is far from perfect.

In fact, there are many words—perhaps most—for which it is simply impossible to detect the physical presence of their referent. Huge numbers of words, of all categories, refer to non-observable entities. That is, a large piece of anyone's vocabulary actually consists of words for *abstract* concepts, which have no concrete physical manifestation in the actual world. How could anyone learn the meaning of a word like *tomorrow* by straight association? There's never any *tomorrows* physically present for the speaker to gaze at, or for the hearer to automatically associate with the sound sequence /təˈmɑɹow/.

The sound-referent association problem arises for even very simple and concrete verb meanings, like *sleep* or *eat* or *wash*. Observation of adult-child interaction shows that in fact most verbs in child-directed speech are used when the event referred to is *not* happening. Utterances like *Go to sleep!* or *What would you like to eat?* or *Let's wash your face* are almost never made while the relevant activity is going on. If the child were to use straightforward association to deduce the meanings of these words, they might think that *sleep* referred to the action of lying down in bed, or that *eat* referred to the action of opening the refrigerator. In fact, they almost never make mistakes like this. How are they learning verb meanings?

7.3.1 Syntactic frames, theta roles and event structure

As noted above, words are almost never uttered in isolation. They nearly always appear surrounded by the functional apparatus that makes up a complete proposition: a declaration, a question, or a command. Even when a word appears as an answer to a question, it nearly always has some functional material attached to it. If older brother says 'What's that?' as the rabbit hops by, Mom isn't likely to say 'Rabbit' as her entire answer. Rather, she'll say 'It's a rabbit', or just 'A rabbit'. At least one of the functional listemes associated with count nouns *has* to appear with any utterance of the content word *rabbit*—it really can't be said on its own. Similarly when answering a question like 'What's it doing?', the answer will be 'It's hopping' or 'Hopping' — not just 'hop'. The suffix *–ing* will invariably appear along with the root verb. As noted above, the only content words that can grammatically appear in an utterance completely free of overt functional listemes are adjectives (Q: 'What color is it?' A: 'Red.'), proper nouns (Q: 'Who is it?' A: 'Peter.') and mass nouns (Q: 'What's that?' A: 'Grass.')

The mere appearance of surrounding structure with a content word can be a powerful clue to the child about what general class of content words the utterance belongs to — mass noun, verb, adjective, etc. — and after learning just a few examples, the child can begin to make generalizations about the likely kinds of meanings that go with items of each class. For instance, *-ing* only attaches to stems that are verbs. If aa child hears a new word with *-ing* attached to it, he will confidently sort it into 'the class of words that can have *-ing* attached to them.' Once the child learns even just one or two verb meanings, he can begin to make guesses about what kinds of meanings words in the *-ing* class can have. If the child has already concluded that words of this particular class don't refer to concrete objects, then he can narrow down the possible meanings for any new word with *-ing* on it, completely independently of whatever is going on around him at the time. He can be pretty sure that whatever the new word's meaning is, it'll have to be a 'property' or an 'event' or some other abstract notion.

If the child has just a little bit of knowledge about the kind of meaning we discussed in the last chapter — *argument structure* and *event*

structure — he can restrict the search space of possible meanings for novel verbs even more, again based just on the grammatical context, essentially independently of any actual observations about the real world at the time of the utterance with the novel verb in it.

Remember that verbs tend to fall into general classes according to how many arguments they have, and what kinds of semantic roles those arguments bear. If a child can sort out noun-words from verb-words based on their co-occurrence with certain function morphemes, he could begin to assign verbs to different classes depending on how many nouns the verbs tended to co-occur with.

For instance, consider the sentences with mystery words in them listed below:

(116) a. The blah will fimble the floop.
b. The gau lammaned the pon the rall.
c. He pangled that she fawed.
d. The windle pates copan.

Assuming that a child can sort the content words into classes according to the function listemes they co-occur with, these sentences fall into the following kinds of patterns:

(117) a. The X_A will X_B the X_A.
b. The X_A $X_B$ed the X_A the X_A
c. He $X_B$ed that she $X_B$ed.
d. A X_A X_Bs X_C.

Now, suppose the child has learned a few basic nouns already, and has noticed that they are members of class A, according to the function elements that they can co-occur with. Based on the nouns he knows already, he knows that words of class A can refer to concrete things in the world—objects, animals, people, etc. So even though he doesn't have the foggiest idea of what *blah, floop, gau, pon, rall,* or *windle* actually mean, he could guess that they might have concrete-item meanings.

Then he's got to figure out what *fingle, lamman, pangle* and *pate* could mean. Assuming that the sentences are conveying complete

propositions, he can figure out that *fingle* must have a meaning that can relate two concrete items to each other. Similarly, *lamman* must have a meaning that could relate three concrete items together. The child has figured out that *fingle, lamman, pangle* and *pate* must have relational meanings, and he's figured out how many arguments each of these relations takes. That is, he knows that *fingle* can have a meaning like *hit, touch* or *cook* but not like *sleep* or *laugh*; he knows that *lamman* can have a meaning like *give, send* or *bring*, but not like *fall, pat* or *sit*; and he knows that *pate* can have a meaning like *run, like* or *drink*, but not *put* or *think.*

These discoveries may not seem like much, but they do reduce the search space for possible relational meanings enormously. In combination with a few additional bits of knowledge, children could narrow the search space even further. Psycholinguists are in the process of discovering that it seems very likely that children do have and use the additional bits of knowledge required.

7.3.2 Agent-Patient Protoroles

For instance, if kids generally assume that S-V-O sentences fall into a semantic pattern of Agent-Verb-Patient (rather than Patient-Verb-Agent), they'll have a leg up in figuring out the meaning of the verb. To take a concrete example, if the child hears *Mary fingled the dog*, in a situation where Mary's patting the dog and the dog's licking Mary, the assumption that the subject Mary is the Agent of *fingle*, rather than the Patient, will allow the child to zero in on a meaning for *fingle* like PAT, rather than LICK. In the same scenario, if the child had heard *The dog fingled Mary*, assuming that the dog is the Agent would cause the child to zero in on a meaning for *fingle* like LICK.

Indeed, there aren't many transitive verbs in English (or any language) where the Subject = Agent, Object = Patient assumption will lead the child to a wrong meaning. The main class of exceptions to this are the stative predicates—usually, predicates that describe psychological states. Consider the following sentences:

(118) a. Mary feared the dog.
b. Mary knew French
c. Mary liked green beans.

These are verbs whose subject is not an Agent, but an Experiencer. Here, the Agent-Patient assumption will induce the child to make a wrong guess about the meaning of /ˈfijr/ — they'll think that it describes a situation where Mary was doing something to the dog, rather than a situation where she had a certain emotional reaction to them.

In fact, though, there is a way for children to sort out these verbs into a separate class from event-denoting verbs, and again, it has to do with the functional listemes that they can co-occur with.

7.3.3 Functional listemes interacting with content listemes.

Remember from chapter 6 that verbs that denote states interact with tense and aspect marking differently from verbs that denote events. The sentences below illustrate this again:

(119) a. Mary is kissing John. (true present tense)
b. #Mary is liking John.
c. Mary kisses John. (only habitual meaning)
d. Mary likes John. (true present tense)

In particular, while a child is very likely to hear verbs that denote activities in the present tense with progressive aspect (*be*+*-ing*), verbs that denote states like *like, know, want, have,* etc. almost never occur in the progressive aspect. Rather, they occur in the 'true' present tense, with no aspectual marking at all. Assuming that children are keeping track of which functional morphemes verb stems can co-occur with, they'll sort event-denoting verbs and state-denoting verbs out into separate classes. Then, given a limited amount of evidence, they could come to understand that the Agent-Subject/Patient-Object semantic pattern goes with verbs in the Event class, and that a different semantic pattern goes with verbs in the State class.

Similarly, children can sort nouns into two classes, corresponding to mass and count nouns, based on whether they co-occur with the indefinite determiner *a* or *some*, or the plural suffix *-s*. Once that sorting has happened, and the meanings of a couple of examples of each class has been figured out, the child has a good clue that can help narrow down their guesses.

7.3.4 Simple co-occurrence? Or actual composition?

Above, we saw that even without knowing what each function listeme means, children could use them as markers when they're parsing the speech stream. Flanking function listemes mark the boundaries of the content listemes between them. Furthermore, children can remember, for each content listeme, which function listemes it co-occurs with.

We supposed that when a child learns the meanings of one or two content listemes from each class, he could generalize certain properties of those meanings, and use those general properties to guide future guesses about meanings for other content words from the same class.

That's the most associationistic way of looking at these kinds of effects—for that idea to work, the function words don't have to have any meanings at all; they could just be class markers. But what if children know more than just 'this function word occurs with this set of content words'? Presumably, it's something about the *meanings* of the function words that allows them to combine with certain kinds of content words, but not others. For instance, it's *because* count nouns name discrete, bounded entities that they can form a plural with *-s*—they can form a plural because there can be more than one discrete, bounded entity. And it's similarly *because* mass nouns name amorphous, stuff-like entities that they *can't* form a plural: there can't be more than one amorphous stuff. It's the interaction of the *meaning* of the function word with the *meaning* of the content word that causes the restrictions on co-occurrence to appear, and allows children to detect the classes of nouns in the first place.

To see the difference between meaning-based co-occurrence restrictions and purely arbitrary, class marking co-occurrence restrictions, we have to turn to other languages. Romance languages, like French, Spanish or Italian, use different determiners with different classes of

nouns.[98] In French, for instance, *le* goes with 'masculine' nouns, like *crayon* 'pencil' and *chien*, 'dog', while *la* goes with 'feminine' nouns, like *chaussure*, 'shoe' and *souris* 'mouse'. Unlike the mass and count classes in English, however, the feminine and masculine classes of nouns have essentially nothing to do with meaning. Whether a noun is feminine or masculine is entirely arbitrary. A child learning French must learn and remember which determiner goes with which nouns, but in this case the co-occurrence restrictions will not help the child make more accurate guesses about the meaning of a new noun. Indeed, if the child tries to base his meaning guesses on the gender categories indicated by the determiner, he'll be sadly misled. It seems clear that the only co-occurrence restrictions that a child can use to guide his guesses about meanings are those that are the result of semantically significant interactions between function words and content words. In order for *that* to work, the child must also know what the function word means.

The first person to investigate whether children actually *do* use clues from the syntactic context to figure out word meaning was Roger Brown, in 1957. He showed three sets of 3- to 5-year olds a picture of some unrecognizable spaghetti-like stuff being poured into a bowl. He told the children from the first group, 'Point to some blick.' The second group of kids were told 'Point to a blick.' The third group were told 'Point to blicking.' Sure enough, the children from each group, presented with the same picture, formed different ideas about the meaning of 'blick'. The first group thought that 'blick' meant the spaghetti-like stuff, the second group thought that 'blick' meant the bowl, and the third group thought that 'blick' meant was the pouring action. The only thing that could have caused them to develop these different ideas about what 'blick' meant is the meanings of the different function listemes in the sentences they heard the word in; everything else about the context was the same. It seems pretty clear that the children were basing their guesses about *blick* on the meanings of the function words it was combined with.

---

[98] As we'll see in the next chapter, Old English was like this as well, but it is likely that more of you are familiar with one of the Romance languages than with Old English, so I'll use them to illustrate here.

The tricky thing about this idea is that function listemes' meanings are the *most* abstract kind of meanings there are. It's reasonably clear that a child can learn the meaning of a noun referring to a concrete thing, like a rabbit, just by associating the sound and the object, especially if the child makes a few general assumptions about objects and eye gaze and so on. Many animals can learn this kind of association without difficulty (does your dog know what /wɑk/ means? But how could a child learn the meaning of *the*, or *-ing*, or *some* from a word-to-world mapping? There *is* no real-world observable entity, action, relation, or property that corresponds to the meaning of function listemes. No animal has ever learned the different meanings of *the* and *a*, or the progressive aspect *be -ing* and the perfective aspect *have -en*, no matter how frequent they are. Yet kids can achieve subconscious mastery of these listemes as early as two or three years of age.

The only guess that linguists so far have about how children learn the meanings of function listemes is that they're predisposed to look for items expressing those particular kinds of meanings. In other words, certain kinds of grammatical meanings are *innate*. Children come pre-wired to know that whatever language they are exposed to will probably have ways to indicate tense, aspect, number, definiteness, mood, person, case, and other functional meanings. If that's true, then their job is just to figure out which of this limited set of functional meanings their language in fact marks overtly, and to match up the most frequent morphemes they hear in their input with these predetermined functional meanings. This idea is called the *innateness hypothesis*.

7.3.5 Yes, but where do the words come from in the first place?

We now have the beginnings of a grasp on how children can go from an unparsed speech stream to a complete inventory of function and content listemes, each with their own meaning attached. But they have to have the speech stream to work with! The child's job is to extract morphemes from the speech stream and induce what their meanings are. The speech stream, of course, comes from the adults and older children around the child, who have already accomplished this feat for themselves. So the listemes a child identifies in the speech stream will depend entirely

on the listemes produced by adults and older children... whose listemes depended on the speech of *their* community when they were learning language... and of course *that* community's listemes would depend on the listemes of the older people around *them* when they were children... and so on, back to the beginnings of speech itself. But a quick perusal of Shakespeare or Milton, or even watching a movie from the 1930s or 40s, will reveal that some listemes that used to be common have dropped out of the language, and others that are currently very common are absent from the older forms of English. A trip to another English-speaking country, or another part of your own English-speaking country, or even hanging out with a different group of people in your own English-speaking town, will introduce you to new words and cause you to wonder where other words went. Who decides what words stay and what words go? And where do new words come from? We turn to these questions next.

7.4 Word Origins

Despite the previous cartoon, only very rarely does anyone sit down and deliberately coin a new word, and it's even rarer for such a deliberately coined word to actually become common currency among many speakers of the language. Most 'new' words are created by some innovative manipulation of an already existing word or words.

One problem we face in discussing how words appear and disappear in a language is deciding what counts as a genuinely new word. The meaning of a particular word can change gradually over time, until the connection between the original and the modern meaning is so remote that one feels justified in saying it's not the same word anymore. The language has a 'new' word, but it'd be difficult to pinpoint the exact moment when the new one appeared. The word *bully* used to mean 'lover, sweetheart', and gradually came to have its current, almost opposite meaning over a long period of time. Does the language have a new word? Or is it the same word with a new meaning?

Let's say for the purposes of this discussion that a new sound-meaning connection—a new listeme—is what we mean by a new word. If a change in meaning is radical enough, as in the case of *bully*, an old familiar sound sequence can come to count as a new word.

7.4.1 New words by 'mistake': back-formations and folk etymologies

The lesson of the preceding discussion is that word meanings are re-created anew every time a new child learns a language. This learning process, although amazingly accurate for the most part (accurate enough to allow the learner to communicate effectively with her speech community), allows for some slippage. A child can acquire a listeme that was never in the vocabulary of the person he learned to speak from, given the right circumstances.

One common way in which new listemes can get pulled out of a speech stream that never contained them is via the same process that gives us Mondegreens and the game Telephone: mis-parsing a word or phrase. If the mis-parse has legs—if it makes enough sense—it may well enter the language as a new listeme. In this way, words that were originally

monomorphemic, underived roots can come to be multimorphemic, compositionally created words. Similarly, morphemes that were present in at the word's genesis may be combined into a single unit, or divided into multiple units, by a misparse. Morphemes that never existed before can spring into existence as a result.

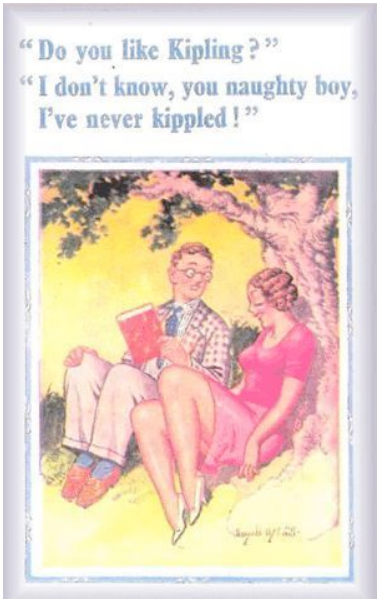

Popular Victorian postcard. The joke is a back-formation.

Misparses come in two main varieties: *back-formations* and *folk etymologies*. The former are so innocuous they hardly seem like a misparse at all. A back-formation occurs when a learner encounters a word that happens to contain a sound sequence that sounds like a particular suffix. The word didn't in fact contain that suffix in the minds of the first people to use it, but the learner doesn't know that. Consequently, the learner's word-analysis machinery strips off the apparent suffix, and induces a meaning for the apparent stem 'stem' by subtracting the apparent suffix's contribution. Some words that entered English that way are *juggle, burgle, televize* and *fluoresce*. The word *burglar*, referring to someone who enters a house to commit *burglary*, was in common use in English in the 1500s. Sometime around 1870, the word *burgle* first

appeared. Where did it come from? Evidently, seduced by pairs like *write-writer, fiddle-fiddler, meddle-meddler,* and *wrangle-wrangler*, some enterprising person who knew the agentive *-er* suffix assumed that *burglar* was made up of a verb *burgle* plus the *-er* suffix. Just as a writer is someone who writes, and a meddler is someone who meddles, a burglar must be someone who 'burgles'. That is, the new verb *burgle*, by the logic of word-formation, must mean whatever it is a burglar does; that is, *burgling* must mean something like 'stealing'. Because the *burgle* + *-er* derivation seems so plausible as analysis for *burglar*, the newly back-formed word caught on, and now is used widely in British English, though not so widely in American English.[99] The same logic underlies the recent introduction of the verb *to lase*, from *laser*, an acronym for ***l**ight **a**mplification by the **s**timulated **e**mission of **r**adiation*.

In this Get Fuzzy cartoon, Bucky is doing something that's more like a folk-etymology. He misparses the sequence /ˈprijməˈdɑnə/ as 'pre-Madonna' rather than the intended 'prima donna', because he's unfamiliar with the latter term. In doing his best to make sense of a string of unfamiliar sounds, he's matched them as closely as possible to listemes that he already has in his vocabulary. This process is the same one that creates mondegreens in song lyrics. The result is called a *folk etymology*

[99] In fact, what Bush did in imagining that *a troop* referred to a singular soldier, from *troops,* was a back-formation. Subtract the plural *-s*, arrive at an unfamiliar stem (to the speaker), and induce a meaning for the stem: if *troops* refers to a group of soldiers, *a troop* must be a soldier.

because it's based on a mistaken idea about the formation of the original word—just like back-formations. Folk etymologies, like back-formations, can give us new listemes; when they do, it's pretty much a judgment call as to whether to call it one or the other. The creation of the word *monokini*, from a root *-kini*, is apparently motivated by a folk-etymology of *bikini*. The two-piece swimsuit originally got its name from a south Pacific island, Bikini, but it came to be analyzed as the prefix *bi-*, meaning 'two' in words like *bicycle, biweekly, bivalent*, etc., plus a stem *–kini*. The stem *–kini* must obviously refer to a swimsuit piece, as *bi-kini* is a two-piece swimsuit; hence half a bikini must be a *monokini*. (One wonders why not a *unikini*, like *unicycle*?) The OED even lists the word *trikini*, built on the same principle.

Similarly, the use of *ology* to mean 'advanced study of some topic' is a folk etymology. In words like *psychology, criminology, immunology*, etc, the root morphemes are *psych-, crimin-, immun-* and *-log(y)*. The *–ology* part isn't a morpheme at all. The *-o-* represents a 'theme vowel', a meaningless vowel epenthesized (inserted) into the middle of Greek and Latin compounds to help them conform to the phonotactic rules of the language: you couldn't have an /kl/ or /nl/ consonant sequence in the middle of words. When a *-logy* compound was created in which the first root ended in a vowel, no additional *-o-* was inserted: *geo-logy, eco-logy.* Two things conspired to make the *-o-* become a part of the *-logy* in the minds of English speakers. From other words with the first part of the compound in them, it was clear that the *-o-* wasn't part of that root: *psych-* appears in *psychiatry*, *psychic*, etc.; *crimin-* appears in *crimin-al, immun-* appears as a word on its own, *immune*. Second, the phonology of English resulted in main stress falling on the antepenultimate syllable in *-logy* words, which nine times out of ten was the 'epenthetic' *-o-* syllable. True epenthesized vowels in English, when they appear, *never* get main stress, so the inclination of English speakers was to assume that the *-o-* syllable must be part of some meaningful root. Consequently all these words came to be analyzed as *psych+ology, immun+ology*, etc, and *-ology* was understood to be a bound root meaning 'study of X'. Since /ˈɑlədzij/ is a perfectly well-formed English word, phonotactically, the step from bound to free was a short one, and now *ology* is an independent word of English.

Back-formations can also result in new bound roots entering the language as well. Indeed, there was a stage when *-ology* was itself a novel bound root. Several new bound roots of this type have entered the language in jocular or slang contexts recently: *-licious*, as in 'bootylicious', which is back-formed from *delicious*, is one example; others are *–tacular* from *spectacular*, as in 'grosstacular', and *–adelic* from *psychadelic*, as in 'shagadelic'. Older bound roots from back-formations or folk etymologies include *-arama* from *panorama*, *-eteria* from *cafeteria*, and *-athon* from *Marathon*.

7.4.2 New words by economizing: clippings

A somewhat related way in which new words can be formed is *clipping*. In clippings, a multisyllabic word is reduced in size usually to one or two syllables. It's often the case that a word is clipped because it comes into more common usage — its frequency count increases — and speakers find that they don't need to use the full sesquipedalian version to identify the concept. Rather, they prefer a more quickly and easily pronounced version. Some words that have come into the language this way are *fridge* (from *refrigerator*), *fan* (from *fanatic*), *mike* (from *microphone*), *fax* (from *facsimile*), *ammo* (from *ammunition*), *flu* (from *influenza*) and many more. Some of these may also qualify as back-formations: *refrigerator* contains several sub-parts that look like English morphemes: *re-*, *-er*, *-at(e)* and *-or*; perhaps the *fridge* part is just the part that speakers assume is the stem of the word. *Burger* is clipping from *hamburger*, but it probably also qualifies as a folk-etymology, since *ham*

is a meat-related morpheme of English, although that word isn't a morpheme in the original word *hamburger*; burgers are made of beef, not ham, and the original item was named after the German town Hamburg. Indeed, *burger* often occurs in a compound with other words that specify what kind of burger it is: *buffalo burger, veggie burger, beefburger;* or what's on it: *cheeseburger* — which perhaps supports the idea that *burger* itself is understood to refer to a generic fried-patty sandwich, and a compounded element is required to indicate what the content of the sandwich is.

Most clippings follow specific, phonologically determined patterns, though, and don't pay attention to morpheme boundaries. Clippings tend to retain the syllable of the word that bears main stress, or occasionally secondary stress. If the clipped version has more than one syllable, the stress pattern is nearly always trochaic, like most common English words, rather than iambic. In going from *raccoon* to *coon*, or *opossum* to *possum*, the initial unstressed syllable is dropped.

Nicknames are created by clipping: *Pete* from *Peter*, *Sue* from *Susan*, *Jeff* from *Jeffrey*, etc. My own name, *Heidi*, was originally formed by clipping an old German name*, Adelheid* and adding the Germanic diminutive suffix *-i,* used in many modern English nicknames. The connection to *Adelheid* is largely lost, now, however, and *Heidi* is usually not recognized to be a nickname at all; it has entered the language as an independent name in its own right. The name *Elizabeth* has both primary and secondary stressed syllables, /əˈlɪzəˌbɛθ/, and the many nicknames formed on this long name use one of these two syllables as the intial stressed syllable in the trochaic ones: *Lizzie, Lizbeth, Liza, Lisa, Betty, Libby* and *Bessie,* as well as *Liz, Bess* and *Bets* are all versions of *Elizabeth*, some of which, like *Heidi*, have entered the language as independent names. The phonology of nicknames and clippings is the topic of considerable study.

7.4.3 *Extreme* economizing: acronyms and abbreviations

A relatively new source of new words and listemes that is becoming increasingly important, especially in English, is a kind of extreme clipping: using the initial letters of the content word in a phrase to

stand in for the whole phrase. This process has been around in English for a relatively long time (*C.O.D.* and *P.D.Q.* originated in the 1800s, e.g.), but it really took off as a new means of word-formation in the second half of the 1900s.

The whole family of inventions is called *initialisms*, and it has two main subgroups: *acronyms*, which are a collection of initials that are pronounced as a single phonological word according to the spelling conventions of English, and *abbreviations*, where the letters are read out one at a time. *AIDS*, /ejdz/, from Acquired Immune Deficiency Syndrome, and *SARS*, /sɑɹz/, from Severe Acute Respiratory Syndrome, are acronyms; *MS*, /ɛmɛs/, from Multiple Sclerosis, and *ASAP*, /ejɛsejpij/, from As Soon As Possible, are abbreviations.

The proliferation of initialisms was a natural outgrowth of a proliferation of bureaucratic institutions named with long, unwieldly compounds and phrases, in particular in the U.S. Army. Franklin Roosevelt's administration (often referred to as FDR, rather appropriately), initiated many programs with such names, which commonly came to be referred to by their initials: the CCC, the WPA, the CWA, PWAP, FERA. These programs and their initialisms are long gone, but the floodgates were opened. The tendency of the American armed forces to initialize everything within sight also had a big impact on common usage around this time, since such a large percentage of the population was involved one way or another with the military during WW2. Initialisms like *G.I., AWOL, SNAFU, radar,* and *sonar* entered the language during this period.

Since then, initialisms have become a completely accepted way of referring to organizations in American English. Very often an organization or group will pick a phrase for their name based entirely on the word that will result when it's initialized. One of the many on-line initialism lists, for instance, is *Ben's Incredible Big List of Initialisms and Acronyms* — BIBLIA for short. It's actually one of the smaller such lists; many acronym dictionaries for institutional and scientific use contain tens of thousands of entries. It's likely that you can think of several local acronyms that are familiar in your own school, workplace or town, but which would be mysterious to anyone outside your own community.

Ben's BIBLIA list is actually a list of a specialized and relatively new type of initialisms: initialisms that have come into use primarily in electronic communications of one kind or another: email, instant messaging, and chatrooms. People who are typing, rather than talking, have a particular impetus to economize, particularly on frequently used phrases, or phrases that are inserted to maximize communicative flow rather than convey actual information. Consequently, a barrage of new initialisms have appeared, some more familiar, some less: IMHO (In My Humble/Honest Opinion), LOL (Laughing Out Loud), MOTOS (Member Of The Opposite Sex), ROTFL (Rolling On The Floor Laughing), RTFM (Read The Fucking Manual), YKWIM (You Know What I Mean), and many, many others, some specialized to a particular group or chatroom (for example, Dave Barry's blog uses OIYDWYMTTY(NY)G to alert readers to certain kinds of links), others a private joke between just two or three people. The main thing of note is that abbreviations and acronyms are now being formed on a daily basis by millions of Internet users. Most will die the day they're coined, but a few will persist, and the net effect will be that new listemes will enter the language.

In one way, initialisms and acronyms are an extreme form of clipping. Especially in military-speak, acronyms and clipping cohabit comfortably in several listemes. Clipped words can be compounded, as in *CENTCOM,* from CENTral COMmand, or a phrase can undergo a combination of clipping and initializing, as in UNSCOM, the United Nations Special COMission, or COMDEX, the COMputer Dealers' EXposition. These initialisms blend clipping and initializing freely, partly to create a final form that will be easily pronounceable as an English word.

7.4.4 Building new words: putting morphemes together

Obviously, in one way the entirety of chapters 4 and 5 were about where new words come from: if you put together two morphemes to express a concept that haven't been put together regularly before, you've created a new word. Bush created a new word in that sense when he combined *security* and *-ize* and *-ation* to make *securitization* with the meaning 'to make secure'.

People make new words like this all the time, of course. Very often, though, the newly formed words don't survive very long. In most cases, someone has 'made a new word' in the sense that they've created a new *phonological* word. If a word is clearly compositional, it's not right to say they've made a truly new *listeme*—it's made up out of existing listemes and their meanings, and the meaning of the whole is entirely predictable. It's only when a new combination of morphemes comes to be slightly idiomatized, not strictly compositional, that we really have a new listeme. At *that* point, it's fair to say a new word has entered the language. *Awesome* and *terrific* used to have compositional meanings when they were first formed — 'inspiring awe' and 'inspiring terror', respectively — but after a while, they became idiomatized: the language had a new listeme, whose meaning wasn't connected to the meanings of its component morphemes.

We've looked at derivational and inflectional morphology in detail already, but there is one process at the cusp of the morphology/syntax interface that we haven't touched on, and which is responsible for many of the listemes in modern English: compounding.

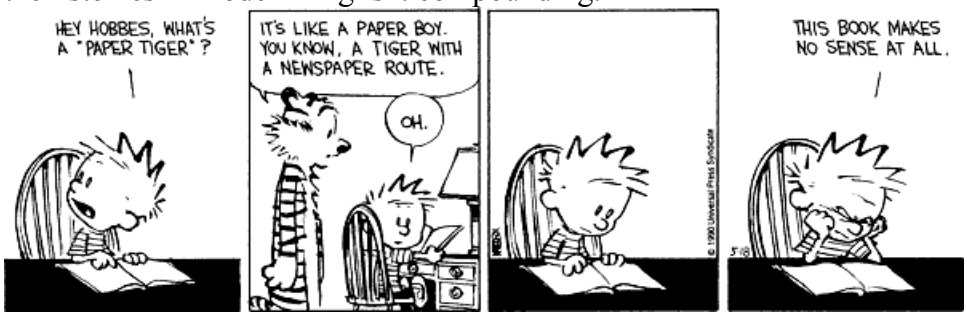

Compounding occurs when two independently meaningful roots are directly combined to form a new, complex member of one of the lexical categories noun or adjective. Examples of each are given below:

(120)

Adjective	Noun
headstrong	high school
skin-deep	rattlesnake
easygoing	sunshine
white-hot	hubcap
outspoken	afterthought

In fact, noun-noun compounding is completely productive in English; new compound phrases are made up every day. If you know what a *scandal*, an *investigation* and a *committee* are, you know what a *scandal investigation committee* is, even if you haven't ever heard the compound before. And if you know *chairperson* (itself a compound), you know what a *scandal investigation committee chairperson* is. And you probably know what a *scandal investigation committee chairperson appointment meeting* would be... and what a *scandal investigation committee chairperson appointment meeting ruckus* would be... and what a *scandal investigation committee chairperson appointment meeting ruckus investigation* would be... and just imagine if they formed a committee to investigate *that*, it would be ...

Compounding creates a new, multiword item that behaves like a member of a lexical category, like N, rather than a phrase, like NP, so it's often considered part of the domain of morphology and word-formation, even though it often involves clearly separate phonological words. (For instance, you can't interpose an adjective into the string of nouns in a compound, unless the adjective is itself compounded with one of the nouns in the string: we have to say *the long$_A$ committee$_N$ meeting$_N$*, not **the committee*$_N$ *long*$_A$ *meeting*$_N$.) And as Calvin discovers in the cartoon above, the intended semantic relationship between two nouns in any given compound can vary unpredictably from compound to compound: *alligator shoes* are shoes made *from* alligators (skin), but *nurse shoes* are not shoes made from nurses, rather, they're shoes made *for* nurses. In order to interpret a compound correctly, speakers often have to understand quite a lot about the way the two elements in the compound are connected, which can often lead quickly to idiomatization. There would be nothing impossible about using the compound *alligator shoes* to refer to shoes made *for* alligators, rather than *from* them, but you'd have a lot of explaining to do to people who have memorized the compound as a listeme, with a fixed meaning.

In any case, compounding with idiomatization, like other kinds of productive morphological processes, is a major source of new listemes. If a compound gets used frequently enough, parts of it may get phonologically reduced, even to the point where it's no longer recognizable as a compound at all. The word *breakfast* is like this.

Exercise 1: Look up the words *lord, hussy,* and *gospel* in the OED or other dictionary with etymological information. What were the root words in each of these former compounds?

So far the examples of compounding I've provided have all been between *free* roots. In fact, some of the most lasting compounds made in modern English are made from bound roots. Scientific and technical vocabulary is nearly all created out of bound roots borrowed from Greek and Latin (we'll discuss why that is in chapter 8). Greek and Latin roots are mostly bound—they can't stand on their own as phonological words—but that doesn't stop them from being independently meaningful listemes in modern English. *Psych-* contributes a similar meaning in *psychology, psychiatry, psychedelic, psychic,* and *psychoanalytic*—and none of these words existed before English speakers created them using the Greek root. Similarly *tele-* is an element of the modern compounds *television, telephone, telekinesis, telegraph,* and *telecast*. Because these roots are bound, it's more difficult to detect exactly what their contribution is in a given compound—you can't look at the root's meaning by itself—but in principle, these listemes are formed according to the same

process of composition+idiomatization that applied to all the other listemes formed by compounding.

7.4.5 *Compounding clips and mixing it up: Blends*

Some new words are created by a sort of combination of all of the above processes. If you clip a couple of words and smoosh them together to make a new word whose meaning is connected to the meanings of the originals, you've made a *blend*. Blends are some of the new words that we're most conscious of, probably because someone had to make them on purpose. Some famous examples are *motel (**mo**tor ho**tel**), smog (**sm**oke* and *fog**), brunch (**br**eakfast* and *l**unch**), chunnel (**ch**annel t**unnel**), napalm (**na**phthenate* and ***palm**itate), guestimate (**gues**s* and *es**timate**)*. Humpty Dumpty described several blends to Alice in *Through the Looking Glass*, though he called them *portmanteau* words: *slithy* from *lithe* and *slimy*, *mimsy* from *flimsy* and *miserable*, and *wabe* from *way before/way behind/way beyond*. The weekly Washington Post Style Invitational often pits readers against one another in creating the funniest new blend just by adding, deleting, or changing one letter of an extant word. Here are some examples from the 2003 contest:

> *Sarchasm*: the gulf between the author of sarcastic wit and the person who doesn't get it.
> *Osteopornosis*: a degenerate disease
> *Beelzebug*: Satan in the form of a mosquito which gets into your bedroom at 3 AM and cannot be cast out.
> *Ignoranus*: A person who's both stupid and an asshole.

Sometimes it's difficult to decide if a word is a blend or a compound of a folk-etymologized root with another morpheme. A good example of this is *infotainment*, which seems like a classic blend, meaning something like **info**rmation enter**tainment**. The morpheme break in *entertainment*, though, is right after an independently recognizable morpheme *enter-* in English, and *info* is a free morpheme on its own now, resulting from a clipping of *information*, so it might be that *infotainment* is really a case of folk-etymology + compounding. It has a sister blend *edutainment*, that might support the independent-morpheme status of

–tainment. In any case, blends show us that people generally feel quite free in manipulating subparts of words to form new words, whether there's morphological justification for the decomposition into subparts or not.

7.4.6 *Meaning change*

If a listeme changes its meaning enough, it can earn itself a new dictionary entry. Our example word *nice* from Chapter 1 is certainly like that. As meanings change, of course, the connection of a word with the root which originally went into it becomes more obscure, which can lead to a pleasant feeling of surprise and discovery when the meaning of the original root is uncovered. Looking at how the meaning of a word changes over time, can give us a little window on how the surrounding culture has changed over time, and hence often enhances historical understanding, not just lexical understanding.

Meaning change is by nature a flexible process, but there are a few recognized paths of change that words can take.

7.4.6.1 *Widening and narrowing*

These patterns of meaning change may be affected by the process of word-learning during acquisition discussed above, particularly the *taxonomic bias* descibed in section 7.2.4 above. Words like *animal* or *retriever* that describe a supercategory or a subcategory of a basic-level word like *dog* are somewhat tricky to learn, because they violate the mutual exclusivity constraint in one direction: anything that can be called a *dog* can also be called an *animal*, though not vice versa. In Chapter 6, we saw that we could organize concepts like these into hierarchical meaning

webs, where connecting lines represented entailing *isa* relationships. Widening and narrowing of meaning can be thought of as a word's meaning moving up and down these taxonomic hierarchies.

A word's meaning *widens* when it formerly used to describe a more specific concept, and over time comes to refer to a more inclusive concept — when it moves up the taxonomic hierarchy. The word *bird*, for instance, used to mean just 'young fowl', but it gradually came to have its broader, modern meaning, which includes all fowl both young and old. Similarly, *manage* used to mean specifially 'handling a horse', but now it means handling anything successfully. Widening often happens as a result of metaphorical or fanciful application of a term. When a learner hears such a use and doesn't understand that it's metaphorical, they simply conclude that the word has a more inclusive basic meaning, that covers a broader range of situations. Children are very literal-minded; mastery of metaphor, humor and meaning extensions is a linguistic skill that is fairly late to develop, and so the metaphorical or humorous nature of a particular usage can easily be lost on them.

Similarly, *narrowing* happens when a word with a formerly broad application is reanalyzed by learners as having a more narrow application—when a word moves *down* the taxonomic hierarchy. Sometimes this happens when another word with a similar meaning comes along and displaces the meaning of the original. This is the case of the word *deer*, which in Old English meant 'animal'. In the Middle English period, though, the French borrowing *beast* came to be commonly used for the meaning 'animal', and *deer* came to be restricted to its current meaning, describing a common kind of wild, herbivorous quadruped. (Later on, the word *animal* was borrowed from Latin, with its modern meaning, and pushed *beast* down the taxonomic hierarchy as well). Other examples of narrowing include a shift in meaning of *accident*, from simply 'happening' to today's meaning, 'unplanned negative event', and *ledger*, which used to be the unmarked word for 'book' but now refers specifically to a book of financial records.

7.4.6.2 Social climbing: *amelioration* and *pejoration*.

Part of the information connected to a word in the mental lexicon is a note about its *register*. Words can be polite, rude, or neutral; suitable for high society or the neighborhood bar. Using a word with a register restriction in the wrong context can lead to negative social consequences, whether it's using a slang term in a formal situation or using a high-falutin' term in a casual situation. A given listeme's annotation for register can change over time, often from high to low, and sometimes from low to high. A word that used to be polite might now be rude; similarly, sometimes a word that used to be casual or slangy might now represent the height of sophistication.

When a word moves from a lower register to a higher register, or from having negative connotations to having positive connotations, we can say that it has undergone *amelioration*. Our example word *nice* from chapter 1 has undergone amelioration; it used to have the negative meaning 'stupid, simple', and now of course it means 'nice'. *Fond* underwent a similar change: in Shakespeare's time it meant 'foolish, crazy, dazed'; over time it came to mean 'over-infatuated, dazed with love'. From there it just came to mean 'in love with' or 'affectionate towards', losing the negative sense entirely.

Pejoration is what happens when a word moves socially downward. We saw that *bully* used to have a positive meaning, 'sweetheart, lover', and we know now that it means 'bully'. It got there via a meaning extension from 'lover' to 'pimp'; from 'pimp', the meaning widened to include not just men who control women's sexual behavior for their own profit, but all stronger people who impose their will on weaker people, particularly for petty reasons. Pejoration is the disease of euphemistic words. The word *retarded* was first applied to developmentally delayed children as a nicer way to describe their condition than the former technical term, borrowed from Greek, *moron*, which had become cruel-sounding through pejoration. Then the same thing happened to *retarded*. A new set of technical terms—*special needs, developmentally disabled* —has begun to replace *retarded* now, again with the intention of allowing a purely medical diagnosis to be made without any implication of social stigma. Changing the word may mask the

problem in a superficial way, temporarily, but it can't help for long without a corresponding change in the underlying cultural attitudes. If there is no such change, then *special needs* will suffer the same fate as *moron* and *retarded* have.

Pejoration is therefore particularly revealing about the underlying attitudes of a given culture. Sociolingusits of English have long noted that terms that were originally neutral ways of referring to the female equivalents of male roles or entities acquired a negative spin that their masculine counterparts lacked. Consider the pairs *mistress/master, spinster/bachelor, bitch/dog,* and *princess/prince*. In each case, there's at least one use—sometimes the most common use—of the feminine term that has negative overtones that the masculine term lacks. The feminine term has undergone pejoration, while the masculine term hasn't. It's argued that this was symptomatic of society's underlying negative attitude towards women: negative attitude, negative connotations. As attitudes change with the change in the status of women in English-speaking countries over the last 50 years, we may see fewer such pairs in the language.

7.4.7 But are these words really *new*?

There are a few other sources of new words in the language. One is the adoption of a proper name as a common noun describing something, often an invention or discovery of the person whose name it is, or sometimes a characteristic way of behaving or speaking that is associated with the person named. Some words derived in this way are *galvanize*, *watt, mach (speed), teddy bear, Kafkaesque, Darwinian* and *sandwich*.

Such listemes are genuinely new, in that the meaning that they get is entirely unrelated to the meaning the root had before, in its use as a proper noun. Most of the other cases we've looked at have involved a manipulation of the form or meaning of an existing word, while still retaining some traces of its original meaning. Still, though, we haven't really seen any cases where the word has really been created out of nothing. Doesn't anybody actually sit down and make up a new words, ever?

In fact, the answer is generally 'No'. People don't make up new words deliberately very often. There are only a few real cases of words being made up out of whole cloth, generally commercial. For instance, the inventor of a new photographic process, Mr. Eastman, invented *Kodak* out of nothing in 1888 to serve as a trade name for his product and company. The word for the number 10^{100}, *googol*, was invented on the spot by a mathematician's nephew, when he asked him what he thought a one with a hundred zeroes should be called. Generally, though, this kind of event is the exception rather than the rule. People get new words from old words.

Throughout this discussion, though, we haven't even touched on the *primary* source of new words in English, which is *borrowing*. The number of new words introduced to English by borrowing makes the combined number of new words added to English via all of the above methods look truly titchy, like a ten next to a googol. To understand borrowing properly, and the remarkable effect it has had on the English vocabulary, we really have to look at the history of English in some detail. We turn to this topic in the next chapter.