

The Monist 84.3 (July 2001) (Special issue The Epidemiology of Ideas, ed. Dan Sperber) 392-416.

Reprinted in The Epidemiology of Ideas, ed. Dan Sperber (Open Court, 2003)

Purposes and Cross-purposes:

On the Evolution of Languages and Language¹

'1. Both the human capacity for language and individual languages have evolved, in part, by natural selection. This paper considers certain aspects and consequences of this, concerning, among other things, the semantics-pragmatics distinction..

Products emerging from histories of natural selection under consistent selection pressures are often described as having "natural purposes," namely, to perform the functions to produce the effects owing to which they were selected. There are many different kinds of selection mechanisms that can generate natural purposes. First, natural selection operates on a number of different kinds of replicators. There are, for example, genes, behaviors established by operant conditioning, behaviors learned by consciously intended trial and error, action alternatives selected through trial and error in thought ("Popperian selection"), a large variety of reproduced cultural items ("memes"²) and so forth. Also, when abilities have been selected for that vary an organism's responses so as to accommodate it to variable features of its environment, these responses have natural purposes. These I have called "derived functions," derived from the original variable ability or "relational function" coupled with a particular environmental circumstance (Millikan 1984 Chapter 2; Millikan 2000). Explicitly represented goals and intentions also represent purposes, of course, although their relation to natural purposes, such as the natural purpose of the kidneys, is more problematic. (More on this momentarily,) In any event, it is clear that there are a good number of different levels and kinds of purposiveness. Theories of language have, I believe, typically collapsed together certain of these levels of function or purpose. In this essay, I will try to keep them distinct, focusing on their interactions in the evolution of language and languages.

'2. Natural purposes or functions, in the sense I intend, result from the operation of selection mechanisms. A natural purpose-making or function-generating mechanism or system is one that has

- (1) The capacity to replicate members of some category of items ("replicators") uniformly, so that copies of copies of copies of these

¹ This paper was largely inspired by Origgi and Sperber, "Evolution, Communication and the Proper Function of Language," P. Carruthers and A. Chamberlain eds., Evolution and the Human Mind: Language, Modularity and Social Cognition (Cambridge: Cambridge University Press) DATE??? Pages???

² Dawkins (1975/89).

replicators are still like the originals in designated respects ("fidelity").³

(2) A tendency to replicate these items selectively depending on their effects, which effects depend, in turn, on respects that are uniformly replicated

(3) Some capacity occasionally to generate novel replicators with the potential to be similarly replicated and selected.

Items that are selected for continued replication by systems of this kind have the effects that accounted for their selection as natural purposes or, as I have called them elsewhere, as "proper functions" (Millikan 1984, chapters 1-2; 1993 chapters 1-2; 2000).

Function-generating systems may select mechanisms with systematically variable capacities or "relational functions," thus producing items, process-stages, activities, behaviors, or whatever, with "derived functions." For example, a heliotropic animal contains a mechanism whose relational function is to cause locomotion in whatever direction the light comes from. When the light comes from direction D, the resulting activity of the animal may be said to have the natural purpose of moving it, specifically, in direction D. Taking a more complex example, baby sea otters seem to be adapted genetically to learn to collect and eat whatever their mothers eat. Their capacities for learning are designed to be focused according to their mother's eating habits. Thus, when grown, different sea otters exhibit behavior patterns with functions or purposes that vary, even though the genetic origins of these purposes are the same. The resulting behaviors of some otters have the purpose of obtaining and extracting the meat of abalones, of others, the meat of sea urchins or of crabs. Most of the specific natural purposes of animal behaviors are, in this way, derived functions.

'3. Purposes derived from natural-function-generating systems, are often contrasted with human purposes. I have argued that this is a mistake, these kinds of purposes not being fundamentally different (Millikan 1984). Human purposes, goals and intentions are merely sophisticated forms of natural purposes. Representations of goals, explicit intentions, and so forth, have accomplishment of their represented ends as derived proper functions. These functions are derived from the complex biological functions of the cognitive and conative mechanisms that have produced them, given the variant inner conditions and outer circumstances of their production. Unfortunately, I can neither explain nor defend this thesis adequately here (see Millikan 1984), but perhaps I can do something to make it a bit more intuitive.

Notice how hard it is to locate the supposed line between merely "natural" and "human" purposes. My plan to buy donuts and milk on the way home from work may seem, clearly, to express a human purpose, whereas the reflex eye blink that prevents a sand grain from entering my eye expresses merely a natural or "biological" purpose, a purpose derived only from a history of natural selection. But is it clear on which side of the line the following purposes belong:

³ "Replicators" and "fidelity" are Dawkins's terms (1976\1989).

The purpose of my blinking after someone, without my awareness, has conditioned my eye blink reflex by smiling whenever I blinked
A subconscious or unconscious purpose, or the purpose of my activity when doing what the hypnotist told me
The purpose with which I gently brake as I negotiate a curve while thinking intently of something else
The purpose with which I automatically push the snooze button on the alarm in the morning, while "still asleep"
The cat's purpose as she meows at the door to be let in
The purpose with which I swing my left foot forward taking a step while walking up the path to my front door.
The purpose with which my cat (or its hind foot) does the same.

Explicit goals and intentions emerge out of a sea of more primitive behavior controls, and the details of the execution of explicit goals are again submerged. The explicit intention to buy donuts emerges from a primitive attraction to sweet tastes, designed to motivate my indulgence in high-calorie foods, although eating high-calorie foods is no part of my explicit intention. The explicit intention to buy milk may have emerged unconsciously from a history of reinforcement by smiles when I drank my milk as a child. And as I retrieve my donut from the package, convey it to my mouth and chew, each minute adjustment of the fingers, hand, tongue, and jaw has a definite purpose, though I am unconscious of most of these motions, let alone of their purposes, even unable to become aware of them explicitly.

Without more defense, I add explicitly represented purposes to the list of levels of natural purpose. Possibly some explicitly represented purposes are represented conceptually (my intention to buy milk and donuts) while others are represented nonconceptually (the purpose of the cat's meowing at the door). We can leave that undecided.

'4. Not all replicators that continue to be reproduced by function-generating systems have functions. Many genes may neither have been selected for nor against, having proliferated by chance (drift), or because they ride close on the chromosome to other genes that do have functions (hitch-hiker genes), or perhaps they are "driving genes," segregation distorters. Similarly, some repeated behaviors may be functionless habits (drift), or derive from mistakes about causes (superstitions). Running barefoot as a child, I accidentally booted our cat running the other direction off the dock into the water. For the rest of her life she ran around that spot on the dock.

In the same way, although the capacity to replicate a great variety of memes undoubtedly itself has functions, there may be many functionless memes. Nervous English speakers say "...uh...uh..." at intervals while nervous German speakers say "...also...also...," nervous Russian speakers produce a series of just audible grunts under their breath and nervous Hungarian speakers, I am told, say "...ö...ö...ö.." Possibly making some kind of filling-in noise serves a function, but

making one rather than another of these particular noises in a given culture probably serves none. Many other conventions, fashions, and so forth also seem to be functionless at least in detail. In the winter it is best to wear something warm, but particular styles in coats, jackets or cloaks are pretty arbitrary. Similarly for many particular aspects of manners and customs. On the other hand, some would argue that there are functions served by conformity to convention or custom simply as such, functions that serve either the individual or the group. It could be, however, that although many details of conventional patterns serve no functions, still conventionality or conformity as such sometimes serves functions. More of this soon.

'5. There are some evolved structures whose purposes are, apparently, to make the evolution of higher levels of purpose possible. Suppose that a function-generating selection mechanism, B, is itself replicated by a prior function-generating mechanism, A. For example, suppose the disposition selectively to imitate the behaviors of one's conspecifics depends on specific genes which are themselves replicators subject to selection. This is no proof that B exists because its selection of products has caused it, in turn, to be selected for by A. Devices without functions may still be replicated. But if any device is improbable enough considering (1) the efficiency and elegance with which it performs certain operations and (2) the complexity of design required for this, and if (3) there is reason to suppose that it probably would have been selected for given opportunity, the argument that it was indeed selected for performing these operations is very strong (compare Dawkins 1982a, 1982b; Pinker 1997,149-175). Thus the mechanisms controlling operant conditioning in the various species surely have making new purposes as a purpose. The fulfillment of these new purposes, derived from learning, are means, in turn, to fulfilling the old purposes of survival and reproduction of the animal or its genes. Similarly, the disposition to play, found in all mammals, may purposefully serve to introduce novelty for the mechanisms of behavior selection to operate on, generating new purposes whose fulfillment fulfills old purposes in new ways. The capacity to learn by explicitly planning and carrying out trials and explicitly noting errors, and the capacity to use imagined actions and results in the selection of plans for action, surely have the making of new useful purposes as their purposes. More generally, the capacities to represent goals, to make plans, to have explicit intentions that can be examined and implemented through explicit inference, surely have all been sculptured by genic selection.

Now consider behavioral "memes," behaviors that are spread by individuals imitating one another. A fascinating question concerns the various mechanisms at work in their selection for reproduction, hence what kinds of functions they can have, and what relation these functions may or may not have to more familiar or fundamental purposes of human behaviors (compare Dennett 1998). A capacity to imitate does seem to have an obvious function in handing down skills from one individual to another. This might happen in either two ways. An individual might observe the outcome of another's activity, and having

grasped what aspects of the activity had contributed to that outcome, might reproduce those aspects to reproduce the outcome. At the other extreme, aspects of another's behavior might be imitated without any understanding of instrumentality. This sort of blind imitation might narrow the range of operant behaviors available for conditioning, thus raising the probability of reinforcement. Useful behaviors proliferated in either of these ways would be selected in the usual way by primary or secondary reward, by furthering explicitly projected goals, and so forth. They would not thus be serving any new kinds of purposes. Also, the mechanisms ensuring a degree of faithfulness in reproduction would be the same as for ordinary trial and error learning, namely, the practical requirement of preserving the instrumental aspects of the reproduced behaviors. Noninstrumental details might be copied too, but these would be vulnerable to copying drift, purposeless migrations in fashion or custom. Skills passed down by copying in this way would not implicate any new kinds of function or purposiveness.

A genuinely new level of selection, hence of purpose, resting on the human ability to imitate has been suggested by Sober and Wilson (1998). Behaviors with an altruistic effect, that is, with a tendency to benefit other members of one's group but at some expense to oneself, are normally at a disadvantage for being passed down genetically. But a general tendency to copy behaviors of others, especially if reinforced by a tendency to demand conformity from others, might result in the selection of groups into which altruistic behaviors had drifted, hence in proliferation of the genes fostering imitation and conformity. Nor would group selection for cultural traits require that the groups selected among be isolated. A higher density of certain altruistic traits in some locales could have this effect.

That is pretty speculative. Altruism to one side, it seems likely that conformity to group practices benefits all just in rendering one another predictable. Appropriate regularities in society are just as important as are regularities in nature for successful planning and action. Especially obvious, behavioral memes make possible the rapid evolution of coordinated behaviors, the paradigm being language use, through which participants in joint activity each benefit at no one's expense. An entirely new level of function for reproduced behaviors then emerges, and it seems certain that our ability to reproduce language forms, at least, has been selected for.

'6. Coordination is required when the contribution required from each participant in a joint venture is determined by the contributions the others will make. Coordination is "blind" when the participants cannot observe one another's contributions prior to making their own. A paradigm of blind coordination is communication with arbitrary signals. What code the signaler should use to convey the message depends on what the interpreter's response will be to the signal, but the interpreter's response should depend on what code the signaler is using. Further, before deciding, neither can directly observe what the other has done or will do. Evolution of primitive signals through genic selection does, of course, occur in many species, but the task for natural

selection is hard due to the lack of strong prior constraints on either side of the signaling process. The generation and selection of learned behaviors is very much more rapid than genic selection, and in the case of human language, at least two kinds of dedicated filters, constraining variety and inhibiting drift in replication, seem to have been genetically selected for. One concerns phonological structures, the other syntactic structures.

There is much evidence that the human auditory systems are specifically designed to accomplish efficient mastery of the phonological structures of a language. Phonological structures are particulate and compositional and they determine what will be heard, in the given language, as the same linguistic pattern repeated again and what as a different pattern. They define the basic same-different scheme for a spoken language, hence what counts as correct reproduction of an element such as a word or a sentence. Thus they help to ensure faithfulness in replication of linguistic representations, enabling the learner to know in advance what aspects of the speech signals produced by himself or others will be the instrumental aspects, the aspects that matter to meaning. Alvin Liberman has argued that phonological structure is the central feature permitting the practical possibility of language innovation (Liberman 1999). Without it we might be stuck with a limited inborn vocabulary, slowly accrued through genetic evolution. Much attention has been paid to the productivity made possible by a grammar that allows embeddings, so that an indefinite number of sentences can be generated with a limited vocabulary. But this kind of productivity would have minimum utility if free to operate only on a tiny vocabulary. The capacity of the language faculty as guardian of phonological structures, thus allowing rapid vocabulary growth not just in the child but also in the public language, makes possible a kind of productivity far more significant.

Universal grammar may be another mechanism helping to effect faithful reproduction of linguistic forms, in this case, syntactic forms. Like grasp of phonological structure, positing universal grammar is positing, at the very least, a filter controlling the aspects of the language a child hears that it will reproduce, or, in practice the same thing, determining what aspects will be perceived as functionally significant aspects. Prior agreement on the kind of materials that are to be used in communication and the aspects of these materials that are to be significant produces a genuinely new kind of faithful replicator, ready for selection.

'7. Besides having their own built in mechanisms for keeping replication faithful, language forms are subject to a new and characteristic kind of selection pressure, promoting the evolution of a genuinely different level of natural purpose or function.

The phonemic and syntactic structures of a language determine which aspects of it will be functionally significant, but do not, of course, determine which significances its specific forms will have. They do not determine what the proper functions or natural purposes cthe meaningsc of various language devices will be. The function of a language device, as with any other product of a

natural-function-generating system, is whatever that device has been doing, whatever effects it has, that account for its continued reproduction. But in the case of language forms (unlike the case, say, of most skills passed down by imitation) which effects will encourage continued replication is not determined by the desires or reinforcement mechanisms only of the agent producing them. The functions of language devices are fulfilled through cooperation between speakers and hearers, hence are determined by the interests of both. Language devices will produce effects that interest speakers often enough to encourage continued replication only if hearers replicate hoped-for cooperative responses often enough. And hearers will continue to replicate intended cooperative responses often enough only if the results are, in turn, of interest to themselves often enough. The function of the language device itself is thus a new sort of function. It is not on the same level as either speaker purposes or hearer purposes. Like items whose natural purposes are derived through group selection (supposing there to be such), language devices have an independent source of function. They are selected for outcomes satisfactory at once to both partners in communication. They have their own natural purposes, often coincident with, but derived separately from, the purposes of individual speakers and individual hearers who use them.

An example will help. Consider, for any language, what gets labeled as an "indicative syntactic form." This form may have a number of alternative functions, just as one's tongue has alternative functions, being designed, for example, to help both with mastication and with speech production. But the form will not be called "indicative" unless one central function is this. It effects production of a true belief having whatever propositional content the various other aspects of the sentence are designed to impart.⁴ This effect is often of interest both to speakers and to hearers. Production of false hearer beliefs may occasionally interest speakers, but rarely serves the interests of hearers. A hearer unable to interpret the indicative sentences he heard so as sometimes to extract genuine information from them would soon cease to form beliefs on their basis. He might first try out other interpretations of the form, and of other linguistic elements used with it, but eventually he would have to give up on it altogether. And if hearers ceased using indicative sentences as guides in forming beliefs, speakers would stop trying to use them to impart beliefs. Production of true beliefs, then, is a linguistic function of the indicative form itself, whether or not a particular speaker and/or hearer have as their purpose to use it that way on a given occasion. Similarly, to instigate actions that accord with their propositional contents is a linguistic function of imperative mood sentences. If it were not sometimes in the interest of hearers to comply with directives, instructions, directions, friendly requests, sanctioned directives, and so forth, imperative syntactic forms would either become obsolete or change their

⁴ This function is, of course, highly relational.

functions.⁵

The replicators that form the basis of a language in use, then, are not utterances only. Cooperative hearer responses are replicators as well. These replicated responses have the same coordinating functions as the utterances that prompt them. Better, the replicators forming the infrastructure of a natural language in use are two-part patterns, an utterance followed by a conventional cooperative hearer response. The whole pattern must be reproduced for the purpose of the language form itself, as opposed to the purposes merely of the current speaker and/or hearer, to be accomplished. A designated kind of perlocutionary act, as well as an illocutionary act, must be performed. Hearers as well as speakers perform conventional acts in using a language for its conventional functions (Millikan 1998).

For any new entrant into an established language community, learning to speak and to understand is acquisition of a purely practical skill. Others in the community are speaking and reacting to speech in habitual settled ways. How the novice must speak and react to achieve the normal fruits of language use merely waits to be learned. Once a specific language is in place, using its conventional forms in the conventional way is what evolutionary biologists call an "evolutionarily stable solution" (ESS) of the coordination problem of communication. No single participant can gain by unilaterally changing the basic rules by which he operates. A settled language community is not easily invaded by incompatible policies of language use. Still, the forms of a specific language are, of course, arbitrary within very broad limits. Languages are thus subject to slow copying drift, slow changes in phonological and syntactic structure as well as in word meaning. And there are also changes that come about through purposive innovation. Speakers and hearers may cooperate to improvise new uses, and these uses may be copied and become conventional. I will discuss these changes soon.

'8 There has been much interest in the idea that memes, like viruses and parasites, might proliferate for reasons of no benefit to their hosts, the people who reproduce them (Blackmore 1999; Dennett 1998). Originally, we must suppose, the general capacity to reproduce cultural forms, or at least the basic ability to imitate, did more good than harm, either to human individuals or to human groups. But this does not imply that the majority of reproduced behaviors either are, or ever were, beneficial. Traits that are only occasionally beneficial are routinely selected for. If group cultural selection has been important, copying things that don't directly aid individuals will of course have been tolerated. Nor if the disposition to imitate has been encouraged through group selection would this imply a disposition to copy only what aids groups. Conceivably some mechanisms determining which cultural forms proliferate easily might be tailored

⁵ For detail concerning this way of analyzing the functions of language forms, see (Millikan 1984).

to produce forms also likely to cause groups to be selected, but there is no more urgency to postulate preselection in this area than for genes. Certainly behaviors that are harmful, even conspicuously harmful, to individuals have proliferated, for example, many fashions in dress, the habit of smoking, and so forth. Useless or harmful memes can indeed proliferate, their memetic functions being merely to do whatever happens, perhaps quite accidentally, to trigger their replication.

Such memes have not simply become independent agents, and this is important to grasp. However, their immediate functions (immediate aims, as it were) do conflict with their own long range functions. A conflict results from the fact that they have functions derived at more than one level, from connected, but still separate natural-function-generating mechanisms. Nor are conflicts of this sort between levels of function something new with the arrival of memes. Consider, for example, the immediate function of behaviors reinforced by saccharine, or behaviors intended to fulfill explicit desires for sweet foods containing saccharine. Contrast these to the long range functions of these behaviors considered as derived from the original biological function of a taste for sweets. Putting this in perspective, notice that the ultimate failure of basic biological purposes is actually more the rule than the exception. The vast majority of individual animals die before reproducing. Similarly, most sperm never reach an ovum, most pounces of the cat yield no mouse, all rabbit chases by foxes end with failure either of the fox's or of the rabbit's behaviors, and so forth. In the same way, derived functions very often do not succeed in implementing the more basic functions from which they are derived. Beaver dam-building behavior is reinforced by anything that will stop the disturbing sound of water trickling, but in the particular case, this something may have no tendency to reinforce a beaver dam. The rat that has been injected with an emetic after eating corn may never eat corn again, even if eating corn becomes necessary for its survival. For my unfortunate cat, running around that spot on the dock had the function (aim) of avoiding sudden inundation, but that was not a function it would ever perform.

Returning to language, two clearly separable sources of function for language forms are, first, the purposes that individual speakers have in using them and, second, the purposes of the forms themselves. If these two sources of function conflict often enough, of course, the result will be either extinction of the form or change of function. Change (or addition) of function occurs if the speaker's purpose is cooperative, the hearer understands the purpose, and this speaker-hearer sequence is reproduced, becoming conventional. Thus the use of Gricean implicature and of various figures of speech can direct the evolution of language in a way that is not merely drifting or random but driven by new speaker purposes that meet with hearer cooperation. Tension between the conventional functions of language forms and the cooperative purposes of individual speaker-hearer pairs that override these functions produces new conventional meanings. To examine these processes in more detail, however, I must first explain the notions of "language convention" and "conventional meaning" that I am using.

'9. As I have been using this notion (see Millikan 1998) "conventions" are merely reproduced patterns of form or activity, where the pattern has a certain arbitrariness in relation to its function (if it has a function). That is, other patterns might have served that function as well: Proliferation of the pattern is due in part to its nature as a precedent. Sometimes a conventional pattern involves one person and sometimes partners or a group of persons. Thus it is conventional to wear white for tennis and black for funerals and to say "Garde la reign" when you put someone's queen in check. Exactly similarly, it is conventional to say "Please pass the butter" when you would like the butter passed and conventional for the butter then to be passed by the hearer. The fact that many conventions serve coordinating functions (see '6 above) is accidental to their nature as conventions, although it is not accidental, of course, that certain kinds of coordination problems can be solved by the emergence or introduction of conventions.

Also accidental to conventions is the fact that following them is sometimes prescribed and sanctioned by social norms of one sort or another, ethical norm, rules of etiquette, laws and regulations, tournament rules, courtroom rules, and so forth. No one frowns if you name your cat "Rover," or if you bake square pies, though naming dogs "Rover" and baking round pies is what is conventional. Nor need conventions be patterns that are in fact followed by all members of some predesignated group. It is conventional for the bride to wear a blue garter, but there is no predesignated group in which this is always done, and sometimes English speakers say "Could you reach the butter?" or "I'd like the butter, please" rather than "Please pass the butter." These latter are all perfectly conventional ways to ask for the butter, illustrating that there are often various alternative conventions available to the same people for doing the same thing.⁶

A pattern becomes conventional, in this sense, if it starts being copied from one person, pair or group to another rather than reinvented each time it appears, and if other equally practical patterns might have been used in its stead and would then have proliferated instead. Conventionality is thus viewed as, in part, a matter of the likelihood of persons doing or saying a thing in a certain way had they not observed its being done or said that way before. Thus conventionality may be a matter of degree, some people copying, say, the use of a certain easy metaphor from friends, while others arrive at it spontaneously.

Notice that this view of the conventions of language contrasts sharply with the view that a language is defined by a set of linguistic rules. Indeed, strictly speaking, on this view there is no such thing as "a language." There are only large raggedy collections of reproducing patterns, many or most of which are familiar to certain ill-defined raggedy groups of people, these conventions functioning pretty well together, mostly not getting in each other's way. And

⁶ For discussion of this notion of convention as applied, especially, to the conventions of language see (Millikan 1998).

because language is just a raggedy collection of reproduced patterns having various origins and independent histories, it is inevitable that linguistic conventions should sometimes accidentally cross. Nothing prevents the same physical sign pattern's emerging from the employment of separate linguistic conventions. Because languages don't rest on prescriptive rules, fully dedicated forms are not possible de jure but only de facto. Glance at the section on equivocation in any informal logic text to find dozens of amusing examples of crossing conventions, the simplest cases involving homonyms. Usually this sort of crossing is harmless. The rest of the sentence or wider context immediately makes it clear from which family of tokens this token has been copied, hence which is its true linguistic function. (The speaker may be using it to serve another piggyback function, of course, but in any event, the hearer will need to grasp the linguistic meaning, the function of the tradition it was copied from, to interpret it.)

' 10. Language conventions, then, are just another variety of "memes," proliferated, usually, because they are serving a cooperative function, and carefully tailored to be replicated faithfully in certain respects. Certain parameters of the reproduction of these memes are constrained by the inborn mechanisms and/or conventions of phonological structure and grammar. But other important aspects of their reproduction are left unfettered, including at least these three:

- (1) The length of the speech stream segment that is copied
- (2) The effects of the segment due to which it is copied, thus what its function or meaning is
- (3) The degree of embedding in extra-linguistic context that may be copied along with it

None of these parameters has a determinate setting for the proliferation of language forms. Consequently, there may be variation in how different persons process the same forms, indeed, in how the same person processes them on different occasions. Thus a clear distinction often cannot be drawn, even in principle, between an expression's having one sense or several senses, between its being used in a different literal sense or only in an extended or figurative sense, between what has been said and what merely conveyed chance, more generally, between semantic and pragmatic phenomena. The attempt to construct intuitive tests or criteria by which these distinctions can always be drawn is, I believe, misguided. They depend at root on statistical regularities in psychological processing, and when different patterns of processing achieve very much the same practical results, there is no pressure for uniformity.

Consider first the length of the reproduced segment of the phonological stream. Just as there is no set length of chromosome that gets copied in sexual reproduction, there is no set length of linguistic expression to be copied. Sometimes a whole sentence form is copied. When answering the phone we say, "This is N" or "N speaking," not "I am N" or "Here is N," but when introducing ourselves in person we say "I am N" or "My name is N," not "This is N" or "N speaking." Other languages do it differently. There are countless conventions of

this sort, people copying one another's phrases rather than conveying the same ideas in other equally possible ways. We speak of a flock of sheep or of geese, a herd of reindeer or of cows, a pack of wolves, a pride of lions, a crowd of people, and once it was a bevy of girls and a blush of boys. We say "next year" for the year after this one and "next week" for the week after this one but not "next day" for the day after this one. Instead, we say "tomorrow," indeed "next day" would not be understood. We speak of the mouth of a river, the mouth of a bottle, the mouth of a balloon, the mouth of a cave, but not the mouth of a house or the mouth of a room, indeed, these last might not be understood at all in normal conversation. Using set phrases rather than composing one's own from smaller linguistic parts, rather, that is, than saying things in "unconventional" ways, is speaking "idiomatically." Beginning at 18 months, children learn about five to nine words a day, steadily, for the next eight years. Their capacity for learning conventional phrasings is probably equally dramatic.

People not only speak in chunks, they understand in chunks. Small children often learn phrases first, taking them apart only later. A close look at the average essay written by a freshman who hasn't understood the material well reveals phrase after phrase that has been swallowed whole and returned still semantically, and sometimes even syntactically, unparsed. These phrases have not been memorized blindly, but understood in a fuzzy, holistic way. Some semblance of their intended meanings has penetrated, but lacking articulation, hence precision. I have myself only recently penetrated "going haywire," "casting aspersions," and "weighing anchor." The meaning of a whole easily separates from the compositional meaning that would be derived from its parts, and may then evolve independently, as in the slippage from "weighing anchor" to "Anchors Away," and from "God be with you" to "Goodby."

'11. Conventions of phrasing are sometimes called "conventions of language use" as opposed to "conventions of meaning,"⁷ but I think this invites confusion. Consider, first, cases in which there has been no slippage so that although the phrase is copied whole, the function performed, the meaning, is the same as the meaning that would be compositionally constructed from its parts. "This is N" said when answering the phone is conventional in the same way that eating with a fork is conventional. Eating with a fork is not made into an eating by the existence of a convention to eat that way. Similarly, given the prior conventions governing the individual components of the phrase "This in N," there is nothing arbitrary about what this phrase means. Meanings of such phrases are readily derivable from living principles governing their parts, even though in practice they may usually be grasped in chunks. Thus the meaning of the chunk is not conventional. But compare a tool, say, a screwdriver, the basic form of which has been reproduced countless times because of its proven effectiveness for driving screws. The fact that it serves this function well has nothing to do with

⁷ Compare (Searle 1975) and (Morgan 1978).

convention, of course. But notice that it has this function as a natural purpose, in the sense defined in '1 and '2, as derived from two different sources at once. First, this purpose is derived from the maker's intention in making it. Second, this purpose is derived from its status as a reproduced item, selected for reproduction on account of prior successes in driving screws. Similarly, where chunking has not resulted in change of meaning, still, it results in meaning simultaneously derived from two sources, (1) derived compositionally and (2) resulting directly from its holistic reproduction to serve that same function. Understanding this makes it easier to grasp what is happening in cases where these two sources of meaning come apart, holistic meaning eventually floating free of compositional meaning.

Where a chunk's meaning has not pulled free from its moorings, some people may hear it as built out of its parts while others hear it holistically. Indeed, whether a given person hears through to its parts may be a matter of degree, or may vary from occasion to occasion. Do these different ways of understanding this string of phonemes constitute it as a strange sort of homonym having two meanings which are, however, somehow also the same? We can get clearer on the matter this way. Linguistic functions or meanings, as distinct from the intentions of individual speakers, are natural purposes that attach to expression lineages. They are functions that account for the continued propagation of these lineages. Homonyms are sound-alikes that form different lineages, so that even if they have a distant common origin, tokens with different linguistic or "semantic" senses have been copied from different branches of the wider family tree. "Table" as in "table of contents" and "table" as in "dinner table" have a distant common origin, but current tokens having one of these senses are never copied from tokens having the other. Like different species, different senses of a word or expression type, listed in dictionaries, for example, as senses "1," "2," "3," and so forth, are distinguished by lineage. Where a phrase is sometimes understood holistically and sometimes taken apart, so long as these different ways of processing don't produce misunderstandings between speakers and hearers, the lineage will remain unbroken through these differences, hence it will have only one sense in the public language.

Similarly, when words or longer expressions are used in extended ways devised by the individual speaker, say, in homemade metaphors, euphemisms or other homemade figures of speech, what makes the usage into an extended rather than a "literal" one is that the expression token was copied from a lineage normally proliferated to serve a different (its "literal") function. Complications enter, however, when a particular extended usage begins to be copied from speaker to speaker, becoming a familiar idiom. It is well known, of course, that in time such an expression may acquire the new use as a new literal sense. But what is happening during the period of transition? What is happening, for example, when a metaphor is dying but is not yet dead? How, exactly, does speciation occur in the case of simple language forms?

What happens is, in part, similar to what happens when phrases are

understood holistically. Different people, whether as speakers or hearers, may hear the expression differently, indeed, the same person may hear it differently on different occasions. Sometimes it carries strong echos from its original usage, sometimes less strong, and sometimes none. Tokens that are either produced or understood with echos are actually hybrids, produced by the crossing of two gradually separating lineages of production and understanding. Tokens that echo for speakers are reproduced by these speakers on two models, copied due to familiarity with the new use but also, in part, due to familiarity with the old. Likewise, tokens that echo for hearers are understood on two models at once. Thus expression tokens at this intermediate stage have two properly linguistic or semantic meanings at once. They continue to owe some of their proliferation to, hence to derive some of their natural purpose from, the function their original lineage still serves, but they also belong to a new lineage with a new function.

Exactly the same analysis applies to half dead conversational implicatures. Tokens of the infamous expression "Can you pass the salt?" no longer function merely as conversational implicatures. They also have a literal imperative meaning. They belong to two lineages at once, a wide-ranging syntactic family of English expressions that proliferates owing to its interrogative function, and a restricted family that is an idiomatic form used for requests. These tokens are both literal questions and a literal requests, nor does this prohibit echoing their origin as conversational implicatures.

In sum, it is not always necessary to make a choice between this being the literal meaning and that being the literal meaning. Occam's razor employed to prohibit proliferation of semantic meanings can be as useless as it is for prohibiting the proliferation of living species or, using a closer analogy, for prohibiting proliferation of addition and multiplication "facts" we memorize rather than re-deriving them each time they are needed. On the other hand, complete separations may eventually occur between lineages that had earlier reproduced tokens jointly.⁸ The classic case results in "The dog went to the bathroom on the living room rug" (Morgan 1978). There is no longer any interbreeding between the families from which tokens of "went to the bathroom" derive.

' 12. Now consider the question, for the sake of reproducing what aspect of its effects is a linguistic segment replicated? That is the same as asking what it's linguistic function, its meaning, is. There is nothing to prevent different people from noticing somewhat different effects, or different effects at different times, or from generalizing to new cases in somewhat different ways. How broad or narrow, abstract or concrete, a category is, how many distinct senses a word has, where literal meanings leave off and extensions begin, are also intrinsically vague affairs.

Consider red hair. If a dress were that color, it would never be called "red."

⁸ Sexual reproduction? The analogies between language form lineages and the lineages of living species are not at all strict. It is better not to get carried away.

Our cat that everyone calls "orange" exactly matches our simulated wood kitchen floor that everyone calls "brown". Do "red" and "orange" each have several senses which are disambiguated according to context? Or are "red hair" and "orange cat" understood as chunks (like "mouth of a river")? Or does "red x" mean red for an x, as "large mouse" means large for a mouse? (Wheeler 1972)? (Is the color that is orange for a cat not orange for a wooden floor?) Does "long" mean the same thing or different things when applied to space and to time, that is, to the first, second, third, and then the fourth dimension? Does "clear" come in different semantic senses when I clear the table, clear the ground, and when the water is clear or the coast is clear? Is the term "lineage" used literally when applied to chains of copied words or is this an extended use? If you say you have "gone around" a squirrel because you have gone from north to east to south to west to north of a tree trunk it is on, even though it too keeps moving around the tree so that you never get behind it, are you extending the meaning of "going around," or does "going around" have two meanings?⁹

Nothing determinate settles questions of this sort, not even within single idiolects. Neither exactly what has been copied from what, nor why it was copied, hence where the copying chains have begun to meander in "new" directions, is anywhere written. Similar remarks go for the traditional question whether such words as "some" and "two" have more than one semantic meaning (at least some, at least two and only some, only two) or whether the shift from one of these meanings to the other is pragmatic. There is no reason to suppose that the various lineages concerned here are clearly formed into separate species, indeed every reason to suppose they are not. Occam's razor, unfortunately, cannot prevent the actual multiplication of complexities.

' 13. Last, consider whether what is copied is expressions or expressions-in-contexts. In the case of language contexts, this question sometimes merges with questions about chunking. Compare "Have you had lunch?" with "Have you had chickenpox?" There is a shift here from meaning "Have you had ... today?" to meaning "Have you had ... ever?" Does this shift take place in the semantics, or is it merely a pragmatic phenomenon? Are there two different semantic meanings of "Have you ... ?" which are disambiguated in context? Or does "Have you ...?" mean the same in both cases, referring vaguely to some past time or other, the more exact time meant by the speaker being pragmatically calculated from context? And there is also a chunking possibility. Perhaps what is copied is the use of "Have you...?" coupled with reference to an event of a sort that happens only once or a few times in a person's life, and another use of "Have you...?" coupled with reference to an event of a sort that occurs periodically. In that case it is not the meaning of "Have you...?" that changes. Rather, the meaning is holistic, not strictly built up out of parts.

Similarly, if John and Bill went to Boston then John went to Boston, so if

⁹ The example is from William James.

Whitehead and Russell wrote Principia why doesn't it follow that Whitehead wrote Principia (Harnish 1976)? Does "A and B" have two meanings, A and B together versus A and B separately, or does it have one vague meaning requiring further pragmatic specification? The chunking possibility is that subject-plus-verb-ascribing-responsibility is copied whole always to mean the subject has sole responsibility. Then the relation between "John and Bill went to Boston" and "Whitehead and Russell wrote Principia" has nothing to do with the meaning of "and." Blind faith in compositional semantics may be rather like belief in "beanbag genetics." Just as the context of other genes in which a particular gene finds itself may radically change its phenotypic expression, similarly, the context of other words in which a particular word finds itself may change the contribution it makes to the semantics of the whole.

There is no reason to suppose that one rather than another of the suggested three ways of copying or understanding how "Have you..." works, or how "X and Y did Z" works, is imposed on speakers of English as such, nor even that it is determinate for individual speakers exactly which aspects of usage they themselves copy. Will it matter to the hearer exactly which aspects of language use are copied? Will it matter which parts are conventional hence semantic and which merely pragmatic? If "Have you...?" has two semantic meanings, then the hearer must use context to determine which of the two is being employed. If "Have you...?" is univocal but indefinite, the hearer must use context to determine, not the semantic meaning, but the speaker's meaning. If "Have you...?" is copied in a chunk along with reference to a periodically recurring event, this chunk having a holistic semantic meaning, understanding still won't be merely decoding, for the hearer will have to determine whether the speaker is interpreting the event as recurring or not. Does "Have you been in Paris?" ask whether you have just been in Paris or whether you have ever been in Paris? How about "Have you eaten snails?" (Suppose the last is asked by a physician examining your hives.) Imagine English-speaking Martians whose nutrition at home comes only in liquid form but some of whom occasionally visit earth. One asks another, "Have you had lunch?", curious about what that strange experience is like.

Where different ways of generalizing are possible, usually it makes no difference at all to the hearer which one is governing the speaker's use. But there can also be contexts in which these differences make a difference. Speakers sometimes purposefully mislead or make jokes by producing tokens that would have one meaning if derived one way but another if derived another.

' 14. Not just position in linguistic context, but position in world context or in conversational or mutual knowledge context, is sometimes copied by speakers. Saying "hit me" specifically when playing blackjack to ask the dealer for another card is copied whole, as is saying "break a leg" specifically to an actor to wish luck in a performance. These phrases-in-context are reproduced items, replicators, each having as-a-whole-including-context a special linguistic

meaning.¹⁰ Contrast the case of "This is N" used to introduce oneself on the phone. Its use in a particular context is copied, but saying "This is N" on the phone does not have an independent meaning. It does not mean what it means because it is said on the phone.

We must be careful here, however. The fact that the whole configuration, "Hit me"-said-while-playing-blackjack, has a special linguistic meaning does not entail that every instance of someone saying "Hit me" while playing blackjack is automatically a replication of this pattern, hence automatically means (literally, semantically) deal me another card. That a certain lineage composed of tokens of a certain sound in a certain context has a certain semantic function does not entail that all tokens of that sound type in that context type are members of that lineage. It does not entail that that configuration of sound and context is somehow dedicated to that function. Indeed, in very principle there can be no strictly dedicated forms in a language it follows that there can be no mere "decoding" of language. True, circles can be drawn around groups of people who then become subject to certain laws, rules, or regulations, including, perhaps, that they are required to use certain designated forms only for certain purposes. A state law may require that Roberts Rules of Order be followed in certain meetings, including, for example, that in these meetings one must not raise one's hand during a hand vote in order to request to speak. In that circumstance, hand-raising is to be a dedicated gesture constituting a vote and it will always count as such. It counts as such de jure. But there are no such circumscribed groups or prescriptions associated with a natural language. No idiom is dedicated de jure, but only, perhaps, de facto, and then only within a certain tradition. Other traditions may always cross over and intercede. Wider context may always suddenly be needed for interpretation.

Demonstratives and indexicals are the most obvious examples of language forms with meanings that depend on context. They also illustrate well how the conventional shades into the pragmatic. Relations of demonstratives to items to which they refer in the immediate external environment, or in the discourse environment, or in the environment of mutual knowledge, may either be copied or improvised. Use of a pointing gesture along with a demonstrative may typically be copied, hence may have a straightforward semantic meaning. But much evidence shows that the use and comprehension of pointing just by itself to focus joint attention is natural for humans. This meaning is not conventional.

¹⁰ True, the context itself may not be reproduced by the speaker. What is reproduced is a relation between a language form and a context. Nothing can be reproduced without the use of any raw materials. The photocopier uses white paper to make a black and white reproduction; it does not make the paper or make it white but leaves it white. Similarly, replicating a relation between a language device and an extra-linguistic context does not require reproducing the context itself. On reproduction, see (Millikan 1984, Chapter 1).

Pointing may still have a semantic meaning derived at another level, however, the way a tool or technique that has been selected for reproduction over and over because it serves a certain purpose especially well acquires a function of its own, a "proper function," no longer directly dependent on the immediate intentions of its maker or user. Much pointing has also become stereotyped or "conventionalized," the particular culture requiring use, say, of the first or the middle finger, or the head, or the lips.

Apart from pointing, there are many other ways to establish joint attention, making salient or using the natural salience, of objects in the external or discourse environment as one refers with demonstratives. Some of these ways are improvised, perhaps, and others copied, one person improvising where another one copies. What is at first merely natural may acquire a semantic function as it is first imitated and perhaps later stereotyped or "conventionalized." Surely this must happen in degrees. To what degree a certain method of demonstrative reference is part of the language and to what degree it is improvised from natural materials hence merely pragmatic will be largely a matter of statistics on individual psychological processing. As such it is not open to a priori inspection or argument.

The domains of quantifiers are understood by their contexts, again, either context in the external environment, or in the discourse environment, or in the environment of mutual knowledge. They refer to whatever domain is currently the object of joint attention, either naturally so, or because the speaker has done or said something to make it so. Possessive forms work similarly. They can refer to any salient relation uniquely pairing possessors with possessed, such as ownership, physical possession, current responsibility for, current use of, an individuating kinship relation, the relation of having been made by or acted on by, and so forth. Thus "John's book" may be the one he owns or carries, the one he wrote or bought or brought, and so forth. Again, context produces the joint attention necessary for communication. In order to be understood in such cases and there are dozens like them the speaker must either recognize or establish joint attention.

Yet the hearer is equally responsible in communication and, indeed, the capacity to recognize exactly where another's attention is directed is extremely keen, even in very young children. The speaker herself, along with her own natural dispositions to attention and the natural symptoms she shows of attention, are a major part of the context. What parts of a speaker's repertoire of methods for establishing joint attention are natural, what parts improvised from natural materials, what parts imitated, and what parts have become stereotyped and conventionalized, cannot be a very determinate matter. Nor can it be determinate exactly how large the copied chunks are, hence exactly what chunks have semantic meaning. Once again, it cannot be determinate where the semantics ends and pragmatics begins. The distinction between the functions or purposes of language forms themselves and the intentions or purposes merely of individual users is again blurred. When misunderstandings occur, often there is

nothing to determine whether the speaker has said or conventionally indicated something wrong, or merely meant something his hearer failed to interpret.

'15. I have sketched some ways in which language is entangled with purposes of various kinds. What I haven't discussed but that is equally important in this context is the forms that speaker purposes and hearer understandings take in the use of language. I have mentioned that human purposes, considered generally, take a number of forms other than the form of explicit intentions, and I would argue that speaker purposes in producing intentional attitudes in hearers are seldom of the explicit kind. I would also argue that in understanding what speakers mean, hearers seldom think about speaker intentions. Obviously the hearer has to understand whatever it is the purpose of the speaker that they should understand, for example, they must understand what they are to do or what information is being offered, but that is quite different from thinking directly about speaker intentions. "The hearer understands what the speaker intends to communicate" is true as a rule only if read fully transparently rather than opaquely. But these arguments will have to wait for another occasion.

References

- Blackmore, S. J. 1999 The Meme Machine (Oxford: Oxford University Press).
- Dawkins, R. 1976, revised 1989 The Selfish Gene (Oxford: Oxford University Press).
- Dawkins, R, 1982a "Constraints on Perfection." In The Extended Phenotype (Oxford: Oxford University Press) 30-54.
- Davis, S. (ed.) 1991 Pragmatics; A Reader (Oxford: Oxford University Press)
- Dawkins, R. 1982b "Universal Darwinism." In D. S. Bendal, ed., Evolution from Molecules to Men (Cambridge: Cambridge University Press) 403-425.
- Dennett, D. C. 1998 "Memes, Myths, Misunderstandings and Misgivings," manuscript presented at The Chapel Hill Colloquium, October 1998.
- Harnish, R. M. 1976 "Logical Form and Implicature" In T. D. Bever et al eds., An Integrated Theory of Linguistic Ability (Harper & Row). Reprinted in (Davis 1991) 316-54.
- Lieberman, A. M. 1999 "The Reading Researcher and the Reading Teacher Need the Right Theory of Speech" Scientific Studies of Reading 3, 95-111.
- Millikan, R. G. 1984 Language, Thought, and Other Biological Categories. Cambridge MA: MIT Press.
- Millikan, R.G. 1998 "Language Conventions Made Simple" The Journal of Philosophy 95.4, 161-180.
- Millikan, R. G. 2000. "Biofunctions: Two Paradigms" In A. Ariew, M. Perlman and R. Cummins, eds. Functions in Philosophy of Biology and Philosophy of Psychology (Oxford: Oxford University Press).
- Morgan, J. L. 1978 "Two Types of Convention in Indirect Speech Acts." In P. Cole, ed., Syntax and Semantics, vol. 9: Pragmatics (Academic Press) 261-280. Reprinted in (Davis 1991) 242-53.
- Pinker, S. 1997 How the Mind Works (New York and London: W. W. Norton & Co).
- Searle, J. 1975 "Indirect Speech Acts." In P. Cole, ed., Syntax and Semantics, vol.

9: Pragmatics (Academic Press) 59-82. Reprinted in (Davis 1991) 265-77.
Sober E. and D. S. Wilson 1998 Unto Others (Cambridge: Cambridge University Press).
Wheeler, S. 1972 "Attributives and Their Modifiers" Nous 6.4, 310-34.