

OXFORD

# Language: A Biological Model

RUTH GARRETT MILLIKAN



# Language: A Biological Model

---

*This page intentionally left blank*

# Language

---

## *A Biological Model*

RUTH GARRETT MILLIKAN

CLARENDON PRESS · OXFORD

# OXFORD

UNIVERSITY PRESS

Great Clarendon Street, Oxford OX2 6DP

Oxford University Press is a department of the University of Oxford.  
It furthers the University's objective of excellence in research, scholarship,  
and education by publishing worldwide in

Oxford New York

Auckland Cape Town Dar es Salaam Hong Kong Karachi  
Kuala Lumpur Madrid Melbourne Mexico City Nairobi  
New Delhi Shanghai Taipei Toronto

With offices in

Argentina Austria Brazil Chile Czech Republic France Greece  
Guatemala Hungary Italy Japan Poland Portugal Singapore  
South Korea Switzerland Thailand Turkey Ukraine Vietnam

Oxford is a registered trade mark of Oxford University Press  
in the UK and in certain other countries

Published in the United States  
by Oxford University Press Inc., New York

© in this volume Ruth Garrett Millikan 2005

The moral rights of the author have been asserted  
Database right Oxford University Press (maker)

First published 2005

All rights reserved. No part of this publication may be reproduced,  
stored in a retrieval system, or transmitted, in any form or by any means,  
without the prior permission in writing of Oxford University Press,  
or as expressly permitted by law, or under terms agreed with the appropriate  
reprographics rights organization. Enquiries concerning reproduction  
outside the scope of the above should be sent to the Rights Department,  
Oxford University Press, at the address above

You must not circulate this book in any other binding or cover  
and you must impose the same condition on any acquirer

British Library Cataloguing in Publication Data  
Data available

Library of Congress Cataloging in Publication Data  
Data available

Typeset by Newgen Imaging Systems (P) Ltd., Chennai, India  
Printed in Great Britain  
on acid-free paper by  
Biddles Ltd., King's Lynn, Norfolk

ISBN 0-19-928476-8 978-0-19-928476-4  
ISBN 0-19-928477-6 (Pbk.) 978-0-19-928477-1 (Pbk.)

1 3 5 7 9 10 8 6 4 2

## *Preface*

Guiding the work of most linguists and philosophers of language today is the assumption that language is governed by rules. Linguists debate over the correct description of the rules of grammar for the various languages, and over the rules of semantics, phonology, morphology, and prosody. And they debate about whether all of these rules are conventional or about which are innately known. A major philosophical literature on the nature of rule-following has amassed largely due to the interest in linguistic rules. And because linguistic rules are taken to be at least in large part conventional, another literature has grown up on the nature of conventions. Those devoting attention to the various uses of language debate what the rules are for the performance of various speech acts and whether all speech acts are conventional. They debate what tests will distinguish what has been communicated conventionally, by semantic rule, from what has been communicated only pragmatically. Many believe that it is of the essence of thought itself to follow rules, rules of inference determining the intentional contents of our concepts, and that these rules originate as internalized rules of language. But throughout these discussions and debates exactly what it is for there to be such things as rules of language remains distressingly unclear.

That the rules of language are not mere uniformities of use follows from the frequency with which they are broken. People lie, they misspeak, they use broken syntax, they misuse words, and they often willfully flout the purported rules with fanciful figures of speech, sarcasm, joking, and the like. So, it is thought, the rules must instead be only 'normative', 'fraught with ought', as Wilfrid Sellars liked to say. But the origin of this normativity remains obscure. From what source do these norms flow? What sanctions enforce them? What happens, exactly, if you don't follow the rules? There is overwhelming evidence that, unlike moral rules, rules of etiquette, rules of the road, rules of games, and so forth, nobody teaches children the rules of language. How then do children learn the rules; for example, the conceptual rules or the rules of pragmatics?

This volume presents a different way of viewing the partial regularities that language displays, the way they express norms and conventions. By ‘normative’ philosophers typically have meant something prescriptive or evaluative, but there are other kinds of norms as well. There are non-evaluative measures from which the facts or from which instances can depart; for example, a simple average is also a kind of norm. I argue that the central norms applying to language are nonevaluative. They are much like those norms of function and behavior that account for the survival and proliferation of biological species. Broadly speaking, they are biological norms. Specific linguistic forms survive and are reproduced together with cooperative hearer responses because often enough these patterns of production and response benefit both speakers and hearers. Like conformity to other biological norms, conformity to these patterns need not be universal or even average. In some cases conformity may not even be particularly common. Conformity is needed only in a critical mass of cases, enough to insure that the cooperative use constituting the norm—the convention—continues to be copied, hence continues to characterize some interactions of some speaker–hearer pairs. Similar norms govern the primitive communication systems of animals, though in that case the reproduction of cooperative patterns of interaction is transmitted genetically rather than culturally, or rather than conventionally.

On this view of language it becomes apparent that what needs to be reproduced, often enough, for a given language to survive is not specific conceptual rules, not uniformity in responses to sensory stimulations or uniformity in inference patterns. What must be conserved are only satisfaction-conditions concerning distal objects and properties, and essential elements of hearer responses to the various forms constituting a language. Because there need not be inferential uniformities (common ‘conceptual rules’) exhibited within language communities, the psychological processes that support our uses of proper names, of words for kinds, properties, and so forth, need to be examined anew. The result is a fairly uncompromising rejection of conceptual analysis as a tool in philosophy.

Another result is that the distinction now generally acknowledged between the propositional content and the force of a linguistic utterance comes at last into very sharp focus. Force emerges as essential to the creation of content rather than as something added to content. A fresh view of the notion of illocutionary force is required, and because what is

conventional in language use involves hearer reactions as well as speaker purposes, the traditional distinction between illocutionary act and perlocutionary act must be rethought. The distinction between linguistic meaning and speaker meaning is also illuminated in a new way, as is the semantics/pragmatics distinction, the distinction between what is said and what is meant.

As language is characterized not by rules but (broadly speaking) by biological norms that concern satisfaction-conditions and function, so is thought. I have discussed thought at length in other places (1984, 1993, 2000, 2004). Here I explore the interface between language and thought. On the model proposed, neither the Quine-Sellars-Brandom interpretation of thought as internalized language nor the Gricean interpretation of language as externalized thought is correct. Neither the intentionality of thought nor the intentionality of language is derived from the other. Although closely entwined, they have separate origins. Also, the processes involved in understanding language are not well modeled by Grice's analysis. They are modeled better as a form of direct perception of the world as mediated, for example, through the natural signs contained in the structured light that allows vision. The main arguments for this position appear in my (2000, 2004); here I explain the consequences for pragmatics, for how language is understood by adults, and for how children learn language.

Many of these chapters have had earlier published incarnations, or at least parts of them have. Their origins are acknowledged in footnotes. Some basic themes and examples appear in several chapters. I have let this stand. Many thanks to Karl Stocker who prepared the index.

## REFERENCES

- Millikan, R. G. (1984), *Language, Thought and Other Biological Categories* (Cambridge, Mass.: MIT Press).
- (1993), *White Queen Psychology and Other Essays for Alice* (Cambridge, Mass.: MIT Press).
- (2000), *On Clear and Confused Ideas* (Cambridge: Cambridge University Press).
- (2004), *Varieties of Meaning* (Cambridge, Mass.: MIT Press).



*This page intentionally left blank*

## Contents

1. Language Conventions Made Simple	1
2. In Defense of Public Language	24
3. On Meaning, Meaning, and Meaning	53
4. The Son and the Daughter: On Sellars, Brandom, and Millikan	77
5. The Language–Thought Partnership: A Bird’s Eye View	92
6. Why (Most) Kinds are not Classes	106
7. Cutting Philosophy of Language Down to Size	121
8. Proper Function and Convention in Speech Acts	139
9. Pushmi-pullyu Representations	166
10. Semantics/Pragmatics: (Purposes and Cross-Purposes)	187
<i>Index</i>	221

*This page intentionally left blank*

# 1 Language Conventions Made Simple\*

---

At the start of *Convention: A Philosophical Study* David Lewis says 'It is a platitude that language is ruled by convention', and he proposes to give us 'an analysis of convention in its full generality, including tacit convention not created by agreement' (1969: 1). Almost no clause of Lewis's analysis has withstood the barrage of counterexamples over the years,<sup>1</sup> however, and a glance at the big dictionary suggests why, for there are a dozen different senses listed there. Left unfettered, convention wanders freely from conventional wisdom through conventional medicine, conventions of art and 'conventions of morality' to conventions of bidding in bridge.<sup>2</sup> Surely it is unwise to try to fell these all with a single stone. Lewis's original goal, however, pursued further in 'Languages and Language' (1996: 134–55), was to describe the conventionality of language, and this may be a more reasonable target.

This chapter is aimed at that target. It is aimed, more specifically, at describing the conventionality of natural as opposed, say, to stipulated language. I do not claim that this kind of conventionality is the only

\* First printed in the *Journal of Philosophy*, 95. 4 (Apr. 1998), 161–80.

<sup>1</sup> For a good discussion of the various difficulties see Margaret Gilbert (1989/92).

<sup>2</sup> *Webster's Third International Dictionary*. Senses 3, 4, and 5 alone—those most pertinent to Lewis's aims—include the following variety: '3. Agreement or an agreement; specif., an agreement enforceable in law; a contract; covenant; 4. General agreement or concurrence, as the basis of any custom, opinion or the like, or as embodied in any accepted standard, method, usage or the like; hence, arbitrary or inflexible custom; fixed usage; conventionality; as slaves to conventions; 5. A rule or usage based upon general agreement; a rule or practice generally adhered to; an arbitrary or inflexible rule, form, principle, etc., as in art; a conventionalism; as, the conventions of morality or of everyday life.' It is not surprising that an attempt to sum up these uses under a single 'analysis of convention in its full generality' should fall to counterexamples.

kind of conventionality there is. I do not claim there are no other senses in which even languages, nonnatural ones, might be 'conventional'; and I do not claim that it is definitional of all language that it has to be conventional. (Perhaps there are animal 'languages', languages that are inherited rather than conventional in the sense I shall describe.)

Why should we care about the conventionality of natural language? Because understanding it is essential for understanding many aspects of pragmatics: for example, questions in speech-act theory and about how natural-language demonstratives work. I shall remark in section 7 on the relation of the analysis I shall give to speech-act theory.

Lewis's analysis of conventions was very complex, involving solutions to 'coordination problems' defined in a complex way, regular conformity to convention within a group, mutual knowledge of this conformity supporting its continuation on rational grounds, and so forth. I shall try to capture the conventionality of natural language in very much simpler terms, displaying its continuity with much more rudimentary conventions. I shall call this kind of conventionality *natural conventionality* and speak, correspondingly, of *natural conventions*. Natural conventions will be described as requiring neither coordinations, regular conformity, nor rational underpinnings.<sup>3</sup> Natural conventionality is composed of two, quite simple, related characteristics. First, natural conventions consist of patterns that are 'reproduced' in a sense to be defined. Second, the fact that these patterns proliferate is due partly to weight of precedent, rather than due, for example, to their intrinsically superior capacity to perform certain functions. That is all. I shall discuss these two characteristics in turn, explaining them as they characterize, first, simple noncoordinating conventions, then simple coordination conventions, and, finally, language conventions, most though not all of which are coordinative.<sup>4</sup>

<sup>3</sup> I have inveighed elsewhere (1984: ch. 3) against the psychological reality of Gricean intentions, hence of 'mutual knowledge' in Lewis's sense.

<sup>4</sup> It is doubtful, for example, whether saying 'Damn!' when annoyed and 'Ouch!' when it hurts are coordination conventions. Expletives are frequently uttered by persons when alone, nor can one question an explicative with 'Were you talking to me?' The conventions of saying 'Boo!' to scare someone and of calling cats with 'Kittykittykitty' (USA) or 'Pusspusspuss' (England) or 'Tssttssttsst' (Hungary, Poland) are not coordination conventions. Lewis says that guiding horses with 'gee' and 'haw' is a coordination convention among drivers who must drive the same horses, however, and perhaps he is right.

## 1. *First Requirement: Reproduction*

The various conventional aspects of natural-language forms are reproduced structures. Whole sentences and larger units are not usually reproduced, of course; but words are reproduced and—even if we humans all come with a Chomskian inherited universal grammar—many aspects of syntax are also reproduced. A pattern has been *reproduced* if its form is derived from a previous item or items having, in certain respects the same form, such that had the model(s) been different in these respects the copy would have differed accordingly. A reproduction is never determined by its model in all respects. The photocopy has a white background not because the original did, but because white paper was put in the copier's paper tray. Similarly, if certain features of syntax are determined by the endogenous structure of the language module, this does not prevent other aspects of syntactic patterns from being reproduced. It is by reproducing language forms that children growing up in France learn to speak French, children in China, Chinese.

Being perpetuated by reproduction is the first basic feature of natural conventionality. Natural conventions are handed down. Consider, by contrast, the *shirt-buttoning* way of proliferating a pattern. Suppose that we all button our shirts in order from top to bottom, but that we do so quite independently, each having independently discovered for herself that this is the easiest way to get the right buttons into the right holes. This convergence of behavior would not result from our shirt-buttoning techniques being handed down, hence would not be conventional behavior. Behaviors that are idiosyncratic or peculiar seem quite certain not to have been handed down, hence these are the paradigm nonconventional behaviors—the *unconventional* behaviors.

On the river in Cambridge, England, it is conventional to punt standing on the deck at the back of the boat whereas in Oxford one turns the deck to the front and stands at the back inside. In both places you see many other, more awkward, antics besides, performed both by undergraduates and visiting untutored novices. Most untutored novices, however, either stand on the deck at the back or turn the deck to the front and stand behind and inside. In so doing, they are not *following* the tradition handed down by either university. Rather, their behaviors accidentally

coincide with these traditions. Their behaviors are no more conventional than the behavior of a puppy who accidentally wets in the bathroom. A behavior is conventional not because its form matches a conventional one but because its form was produced by reproduction.

There are a number of different ways in which conventional patterns can be reproduced which will be of interest to us. They may be reproduced by being copied from one another directly, or one person may tell another how a pattern goes. For example, Johnny's mother tells him that he is to put his letter in the mailbox and put up the flag. The result is that he reproduces the first part of a conventional pattern of activity, the second part of which will be reproduced by the mailman, who, on seeing the flag, stops at the box, removes its contents, then puts the flag down. He does this, in part, because he is familiar with this conventional pattern of activity and, recognizing that the first part has been performed, reproduces the second part. With his addition, the full conventional pattern has now been reproduced.

Sometimes conventional patterns are reproduced by what I shall call *counterpart reproduction*. A bolt one wishes to be suitable to a given purpose usually has various thread gauges as reasonably functional alternatives, but once a particular nut has been chosen, the choice of suitable thread gauges is strictly determined. Likewise for the nut. So if certain nuts were made to fit certain bolts, and then more bolts made to fit the nuts, and so on, specific thread gauges of both nuts and bolts would be reproduced without any bolt's being directly copied from any bolt or any nut from any nut. The traditional positions assumed by men and women for ballroom dancing were commonly reproduced in part by *counterpart reproduction*, each woman settling into the traditional woman's posture in response to the postures of the men with whom she danced, and vice versa. A simpler kind of counterpart reproduction is *handshake reproduction*. Shaking hands by placing the right hand in another's right hand, thumb notch to thumb notch, then shaking up and down (contrast left hands, grasping wrists, interlocked fingers, shaking side to side) is probably learned in part not by direct copying but by fitting in with one's partner; but here the counterparts must be alike to fit. Often each individual part needs to fit arbitrary members from a whole collection of counterparts, and vice versa. This easily

results in standardization of forms, more easily than by direct copying. Copies of copies easily drift away from the original; the need to fit counterparts retards drift.

That a certain pattern of behavior is common in one culture but not in another is evidence, but not particularly strong evidence, that the pattern may be being proliferated by reproduction. One reason it is not particularly strong evidence is that the only practical manner of doing a thing may depend upon context and hence vary from culture to culture, proliferated in the shirt-buttoning way. The sensible way to dress in Iceland is not sensible in Ecuador, copying quite to one side. But the reasonable way to behave may also depend upon prior patterns already established in one's culture, so that a pattern learned merely as practical, as a skill, at a deeper level of analysis may be reproduced by counterpart reproduction. Margaret Gilbert says that in some circles it is conventional to write a thank-you note after attending a dinner party but that, *contra* Lewis, there is nothing arbitrary in doing so (1989/92: 340). Indeed, if the host expects you to write such a note and will think you rude or ungrateful if you do not, then you have reason to write it, conventions quite to one side. On the other hand, the host will have such an attitude only because people before you have written thank-you notes under such circumstances. Had people before you made phone calls, or sent messages or flowers, your host would have expected that instead. Thus, even if you were not aware of doing what others have done, writing a note in response to your host's expectations, however you chanced to discern them, is a reproduced behavior; it is handed down. Reproduced patterns are often learned rather in the manner of skills, or because obviously morally required, and so forth, while at a deeper level they are portions of counterpart-reproduced, conventional patterns.

Patterns of activity are often reproduced unconsciously. For example, social distances—the distances at which people stand when conversing—vary from culture to culture and from one kind of social relationship to another within each culture. If you are standing at the wrong social distance, the person to whom you are talking will move; so to prevent slow circling about the room as you talk, you unconsciously reproduce the conventional social distance by handshake reproduction. Shaking hands with the right hand specifically, may sometimes be reproduced in



this unconscious way. One can even imagine being forced, without realizing it, into the right-hand driving pattern merely in an effort to avoid oncoming traffic. Then one would unconsciously reproduce a driving pattern by handshake reproduction.

The reproduction and standardization of natural-language patterns may take place, in part, through unconscious counterpart reproduction. Reproduced speaker–hearer patterns generally involve inner as well as outer acts so that direct copying is not a means of reproduction for the whole. But given the speaker’s portion of the pattern, the hearer’s part is not arbitrary; and given the hearer’s portion, the speaker’s part is not arbitrary. So the ability to play either of these roles may be acquired, in part, as skills, given partners who are reproducing the complementary portions of these patterns. For example, those complex patterns whose standard job it is to transfer information—involving, paradigmatically, the use and interpretation of indicative sentences—are completed as the hearer performs the inner act of interpreting these sentences, in accordance with certain rule-like patterns, into beliefs. This movement is a hidden action and (René Descartes to the contrary) not voluntary.<sup>5</sup> Hearers learn to believe in conformity with the conventional rules of the language because this often brings the reward of useful knowledge gained. Except for the (probably very large) boost from certain inborn language capacities (compare the texture and color of the paper fed into the copying machine), they learn it much as they learn to interpret natural sign patterns. Similarly, speaker–hearer patterns that go with paradigmatic uses of imperatives begin when the speaker expresses a purpose with regard to the hearer’s actions with outer signs in a certain established, that is reproduced, manner. Learning to do this may result, in part, from the speaker observing the effects that his imperatives, and those of others, have on hearers who are already disposed to complete these patterns by supplying conforming actions in the established manner.

<sup>5</sup> Compare Lewis’s claim that forming a belief is not a voluntary action and hence cannot be part of any convention (1969: 180). Later (1975/96), he allows beliefs formed on rational grounds to be part of the ‘content’ of language conventions. But social distances are not adopted in accordance with reasons. Indeed, they are so unconscious that their very existence is known to few people and went unremarked until anthropologists began to study them in the 1950s, yet they are surely paradigms of natural conventions.

## 2. *Second Requirement: Weight of Precedent*

Not all reproduced behaviors are conventional behaviors. Handing down a skill is not, as such, proliferating a convention. I learned from my mother, and she from hers, to open a stuck jar lid by first immersing it in hot water. Opening jars this way is not thereby 'conventional'. To be thought of as conventional, a reproduced pattern must be perceived as proliferated due, in important part, to weight of precedent, not to its intrinsically superior capacity to produce a desired result, or due, say, to ignorance of any alternatives.

Thus, a pattern is considered conventional only if thought to have little tendency to emerge or reemerge in the absence of precedent. This might be because the pattern has no useful function at all: wearing black and white to examinations at Oxford University, decorating for Christmas specifically with red and green, handing out, specifically, cigars after the birth of a boy. Or this might be because there are other viable patterns of behavior that could have served the same functions equally well had precedent for them been established. Thus, a certain arbitrariness of pattern in relation to function—evidence for this perhaps gleaned from other cultures, where other patterns do serve just as well—is an argument for conventionality. Conventional patterns are patterns for which other patterns, given different historical accidents, might as well have been substituted. Lewis, too, emphasizes arbitrariness as essential to conventionality.

Yet an arbitrariness of pattern in relation to function is not sufficient evidence for conventionality. The songs of the various bird species are to a large degree arbitrary in relation to function, but they are not conventional, because they are not copied or reproduced in the sense defined above. That other reproduced patterns would serve just as well if only their possibility was *known* is also not an argument for conventionality. Lighting fires by rubbing two sticks together may be a sort of convention among boy scouts, but it is not merely conventional in a culture that knows no other way to light fires. Similarly, conventional Western medicine is called *conventional* only because other kinds of medicine are now also known in the West. Calling it conventional also implies (perhaps euphemistically) that these other kinds are genuine alternatives,

only weight of tradition serving to exclude them from more common use. But where knowledge of alternatives is unavailable, precedent figures in the proliferation of a pattern not essentially, but only as a means of making its possibility known. If people discovered the pattern's potential independently of one another, it would proliferate just as readily. A conventional pattern, by contrast, proliferates partly due to precedent.

There are various mechanisms through which patterns of activity may proliferate due to weight of precedent. A pattern may prevail over easily invented alternatives merely because it is easier or more natural to copy than to use one's imagination, or because people prefer to do as others do, not wanting to be out of step, or because what is familiar is as such pleasing, or because people feel more secure in the tried and true. In some cultures all traditional behaviors may be positively sanctioned as such, either out of simple respect for tradition, or for fear of disturbing the moral order. It can also be habitual to place sanctions on certain but not other conventional behaviors, for approval and disapproval themselves are often copied without much thought. In all such cases weight of precedent itself is helping to account for the proliferation of the pattern, making the pattern conventional. Sanctions rooted merely in tradition often are not recognized as such but are thought to be independently based in morality, or etiquette, good taste, necessity for health, decorousness, propriety, piety, and so forth. Often there is room for reasonable disagreement over which reproduced behavioral patterns are sanctioned merely conventionally and which are sanctioned due to some function they are serving.

When patterns of behavior are proliferated by counterpart reproduction, often it is weight of precedent which keeps them afloat, and which does so for very practical reasons. In the Orient, because people have learned to use chopsticks, they are the implements placed on the table for dining, manufactured in quantity, available easily and cheaply, and so forth. In part because they are manufactured in quantity, available easily and cheaply, and so forth, chopsticks are placed on the table and people learn how to use them. Similarly for forks in the West. Weight of precedent proliferates their use, and not just because people are

conformists. Proliferation of the use of standard units of measure is another example of this sort. Measurements taken in conventional units proliferate in tandem with the calibrated devices used to make the measurements. For example, measurements taken in yards proliferate as yardsticks proliferate, and vice versa. Weight of precedent proliferates the use of conventional units of measure also by handshake reproduction. Measurements taken in yards are easiest to compare or to calculate with other measurements taken in yards, those taken in meters with others taken in meters, and so on.

### 3. *Coordination Conventions*

Lewis claimed that all conventions solve coordination problems among the participants (1969: 36 ff.). A subclass of natural conventions, I believe, are indeed *coordination conventions*. My description of coordinations will be simpler than Lewis's, however. Especially, no beliefs on the part of the participants will figure as essential ingredients in the analysis.

There is a need for coordination when:

- (1) members of a group ('partners' in a projected coordination) have a purpose in common;
- (2) achieving this purpose requires actions by each of the partners;
- (3) more than one combination of actions will achieve the purpose;
- (4) the set of workable combinations fails fully to determine what any single partner's contribution must be independently of the actions performed by the others.

Coordination is achieved if the partners' combined actions achieve the common purpose. Coordination conventions are conventional patterns of activity that proliferate, in part, (causally) because they achieve coordinations.

Some coordinations require the partners to act the same way, others, to act differently. Coordinating social distance requires acting similarly. Although each person learns independently to maintain the conventional distances (as we may learn independently to button our shirts

from top to bottom), it is not the case that keeping just these distances has intrinsic superiority over keeping other distances, nor that ignorance of other possibilities is a factor in proliferating the distances actually kept. Rather, that others are already doing things this way causes new participants to follow after. Thus, mere weight of precedent perpetuates this coordination solution. Another coordination, discussed by Lewis, which requires similar actions, is a community's using certain standard media (minted money) for exchange of goods. I have mentioned that common units of measure are proliferated in part by handshake reproduction because of the ease of comparing measurements made with a common measure (the 'partners' here might be oneself at different times). As such, the use of standard units of measure is a conventional coordination. In another case discussed by Lewis, different actions are required by the partners to achieve coordination. When a telephone connection is broken, one person must call back while the other waits.

On the other hand, we could say of the telephone partners that they must do the same thing; namely, follow the pattern: original caller calls back, original receiver waits; or perhaps original receiver calls back, original caller waits; or perhaps, whoever's sentence was interrupted calls back, the other waits. If any one such pattern should begin to appear slightly more frequently than others, due perhaps to the vagaries of statistics with small numbers (compare 'evolutionary drift'), granted this is consciously or unconsciously noticed, it will tend to reproduce in a counterparts manner because people learn from experience. Then, like a gene with a fitness advantage, it may 'tend to fixation'. That others are already doing things this way causes new partners to follow after.

Gilbert has argued that there is nothing rational about following precedent in the attempt to solve a coordination problem, unless one makes more assumptions about the nature of one's partners than just that they are rational (1989/92: ch. 6; 1983). What more one needs, I am suggesting, is to assume that one's partners routinely adapt themselves to their environment by repeating what has worked in the past. Coordination conventions proliferate because, rationality aside, people learn from experience exactly as other animals do.

#### 4. *Conventions Do Not Imply Regular Conformity*

There is a long tradition of taking conventions *as such* to involve regular conformity within a group either *de facto* or *de jure*.<sup>6</sup> If conventions are as I have described them, this is a mistake. In this section I shall argue against regularities *de facto*; in section 5 I shall argue against regularities *de jure*.

That people need not regularly conform to noncoordinating conventions is clear. Few actually hand out cigars at the birth of a boy, nor does everyone wear green on St Patrick's Day, or decorate with red and green at Christmas, or punt from the deck when on the River Cam. What, then, of coordinating conventions?

First, I shall argue that to solve a coordination problem it is not usually needed that all sets of cooperating partners in some group solve it in the same way. It will help here to differentiate among three different kinds of coordinations that can be proliferated by weight of precedent. I shall call these *blind coordinations*, *half-blind coordinations*, and *open coordinations*. In blind coordinations each partner must act before having any evidence concerning the other's action(s). Thus, when a telephone connection is interrupted, neither party has information about the other's reaction until the coordination has either succeeded or failed. Contrast two people moving a couch together where each sees the anticipatory motions of the other. Usually, one adopts a leader role, the other a follower role, and an excellent open coordination may be achieved in the absence of any conventions at all. Between these extremes is half-blind coordination: for example, when a child raises her hand to speak in a classroom and is called on by her teacher. Here part of the conventional pattern wishing-to-speak-produces-hand-raising-produces-being-called-on which the child initiates is visible to the

<sup>6</sup> Searle (1969); Lewis (1969; 1975/96); Schiffer (1972); Bach and Harnish (1979); Gilbert (1983; 1989/92); Recanati (1987). Recanati takes it that conventional language devices 'indicate' or, using the linguists' term, 'mark' uses of language; that is, conventions mandate that these devices shall be used only for those purposes. This is not explicitly stated, but see, for example, § 22.

teacher, but the wishing-to-speak part is not. Only conventions for achieving the first kind, the totally blind coordinations, depend for effectiveness on a degree of regular conformity within an interacting group.

Consider the right- (or left-) hand driving convention. The cooperative purpose requires that drivers approaching one another both either drive on the right or drive on the left. There is no secure way for leader–follower roles to be established, the time for decision is small, and it is hugely important that the coordination be achieved. The coordination thus behaves like a fully blind one. The only way to achieve it with complete reliability for every pair is to achieve it always the same way within some group that can easily identify its own members (people driving in America; people driving in England). Regular conformity to the same pattern is required among members of the entire larger group from which the various cooperating pairs emerge.

Contrast with blind coordinations what I shall call *leader–follower coordinations*. Here there is no need for regular conformity. One member leads off in a manner either wholly or partially observable to the other, the other responds so as to reach a mutually desired conclusion. Many such coordinations are completely open, hence completely unproblematic, as when you seat yourself at an arbitrary table in the restaurant and I follow after. The interesting ones, those which are helped by conventions, are those which are partially blind. Consider the couch-moving case more carefully. The follower cannot be guided by the very actions with which he is simultaneously coordinating, but only by their natural precursors. It is natural, then, that the leader should exaggerate and stereotype his anticipatory movements so as to make them easier to follow. Exaggerated and stereotyped anticipatory movements, reproduced by experienced cooperators because they have been previously experienced to produce effective coordinations, are *conventionalized* movements.<sup>7</sup> One more step gives us conventions proper. Let the leader's

<sup>7</sup> The ethologist Niko Tinbergen considers the communicative signals of animals to be 'ritualized' versions of 'intention movements' (1952). An *intention movement* is one from which it is possible to infer what an organism is likely to do next. *Ritualization* is a process whereby specific features of such movements become progressively exaggerated. It is considered likely to occur when the benefits to an organism of conveying information

indication of the action pattern he purposes for himself and his follower be divorced from any natural anticipatory movements. That is, let the reproduced pattern—it begins with a lead from a leader, proceeds through a specified pattern of action involving roles for both leader and follower, and culminates in a certain result—have an arbitrary beginning. Then it is a fully conventional coordination pattern. But there is no need for any whole group regularly to use this pattern for it to be effective in individual cases.

Conventional leader–follower coordinations begin when a leader reproduces a certain portion of a traditional pattern, that portion being observable to a follower. The follower is familiar with the pattern, recognizes it, and reproduces the complementary part, resulting in a coordination of a sort that is partly responsible for the proliferation (due to precedent) of the pattern. When the convention is half-blind, the leader may also have completed an important part of the pattern beforehand but it was unobserved by the follower. Thus, the leader of the mailbox–flag convention places mail in the mailbox before the postman arrives, and the leader of a conventional-language pattern has invisibly made a conventional transition from belief or intention into an outer language. Or perhaps a portion of the leader’s part is completed only later on, as when a promise is made.

To achieve such a coordination, the leader and follower must, of course, be reproducing parts of the *same* conventional pattern. But there is no need for only one pattern to be used in one group. There are an indefinite number of ways to impart any given piece of information using English. It is not even necessary always to speak the same language, even with the same partner. Leader–follower conventions emerge where coordinations are half-blind, but also half-sighted. Enough is seen for the follower to discern with which conventional pattern he is dealing immediately on observing the lead. Following is like chiming in after the first line of a song; there is no need for a community to know just one song. Nor does this sort of convention require that there be some definite group in which it is mutually believed that it is a convention.

about its future behavior exceed the cost of doing so. For instance, conveying information about an intention to fight may be beneficial if it prevents the fight from occurring.



People acquainted with a certain leader–follower coordinating convention need not form a social group, or always know who one another are. It is usually best to know in advance whether I shall be understood by my follower, but the brute existence of a convention does not require this in each case.

Perhaps it will be thought that I have missed the essential regularity that is required by leader–follower conventional coordinations; namely, regular follow-through by both parties in completing conventional patterns once initiated.

Again, consider the case of blind coordinations first. Even there, only regular *enough* observance is required to sustain them. If only one in five friends follows the telephone convention, I shall still be somewhat better off following it myself. Similarly, even if only one in fifty folks that I ask ‘Give me a dime, please, just a dime?’ actually gives me a dime, it may be quite sane to keep at it. On the other hand, if more than a certain critical mass of statements were (nonobvious) lies, our indicative-mood forms might indeed die out. I shall say more about this below.

## 5. *Conventions Without Prescriptive Rules*

Complementing the claim that conventions require regular de facto conformity among members of a group is the common claim that conventions rest on rules requiring conformity *de jure*.<sup>8</sup> In what sense then do conventions involve rules, and what kind of rules are these? It will be well to proceed quite carefully here.

There are many conventional patterns that are usually reproduced by explicitly following a set of rules. This is true, for example, of the conventional patterns constituting most games and of the pattern that is parliamentary procedure. Of course, it is not just the having of rules that makes a game conventional, as is shown when you make up rules for a game, thus playing an unconventional game. Notice also that rules of this sort concern regularities only within conventional patterns, not regularities, either de facto or *de jure*, in their being used. The rules of

<sup>8</sup> Gilbert, for example, explicitly endorses the latter claim but not the former.

chess do not tell you when to play chess, nor do the rules of parliamentary procedure tell you when you must use it. It is just that you are not reproducing that particular conventional form unless you produce a pattern showing that internal regularity: in chess, the king always moves just one square; in parliamentary procedure, the chair never speaks to a motion; and so forth. On the other hand, some conventions are indeed described by rules telling when they are to be applied. There is the conventional rule 'At Christmas, decorate with red and green', and the rule 'When you are the new father of a boy, hand out cigars'. In neither kind of case, however, are conventional rules, just as such, prescriptive. These rules describe conventional patterns; they do not prescribe them. The bare existence of a convention neither mandates, nor gives permission for anything.

On the other hand, conformity to a convention may be or become mandated and sanctioned. Following parliamentary procedure is mandated by law in the United States Congress. Crucially important blind conventional patterns, like driving on the right, are heavily sanctioned (indeed, often written into law), and it is not a mere matter of convention that they are sanctioned. Whenever it is desirable that some predictable social pattern or other be stabilized for some purpose in a society, or that some convention or other be universally followed, then it is not merely conventional that there are sanctions attached to conformity. But there are also conventions that are sanctioned merely conventionally; that is, sanctioned for no reason beyond tradition; and in some societies all or most conventional behavior may be sanctioned simply as such.<sup>9</sup>

<sup>9</sup> The distinction is blurry between (1) patterns that are first merely conventional and only later written into law or otherwise mandated and (2) partially arbitrary patterns that originate only when written into law or otherwise mandated, resulting in a blurriness in what counts as conventional. Compare, for example, the pattern of behaviors that constitutes getting married in a certain state of the United States with the pattern to which one must conform to arrive at United States citizenship. Both are to a degree arbitrary. Neither is nowadays reproduced merely due to weight of precedent, but the marriage ceremony once was. Which, if either, is conventional? If either is conventional, might there be another sense of 'convention' edging in here; say, the kind that, according to *Webster's Third*, is 'formed by agreement'? This sort of blurry line between natural conventions that are sanctioned and what we might call *regulated conventions* may help to produce the illusion that natural conventions are *de jure*.

Just as rules for constructing conventional patterns do not, as such, mandate that you must construct them, they do not, as such, prohibit constructing part or parts of the pattern without constructing the rest. The rules of chess do not mandate that having started a game you cannot then quit in the middle, or that you cannot set up an end game or a middle game and just play that much for the practice. If you cannot quit or change rules in the middle, that will be a matter of external rules or sanctions—say, rules of etiquette, or tournament rules, not the conventions of chess.

Conventions often involve repetitions of complex patterns. But fragments of these patterns may also occur quite freely. Conventional patterns are very frequently broken. This is especially common when construction of a conventional pattern requires the cooperation of more than one person. Consider again, 'Give me a dime, lady, just a dime?' That the pattern, if unbroken, would involve transfer of a dime is clear. The syntactic pattern for directive sentences continues to be reproduced only insofar as the whole speaker–hearer pattern in which it is conventionally embedded is sometimes completed. If nobody ever complied with another's directives, the directive forms in a language would soon cease to be used; indeed, to be understood. Earlier, I made a similar remark about hearers' believing assertives. It does not follow that complying with directives or believing assertives is, as such, mandated by a 'rule of language' understood as prescriptive.

Lewis tells us that the conventions of language are 'truthfulness' and 'trust' (1969:148 ff.; 1975/96:137 ff.), which unpacked means roughly: 'Say only what is true or what you actually want the hearer to do', 'Believe what you hear', and 'Do what you are told'. He is right, I believe, that these rules elliptically describe conventions for certain of our most basic language forms, but he is not right that these rules describe regularities, nor, I am now arguing, do they mandate conformity. Certainly, there are circumstances under which people are under mandate to do what they are told, and under some political regimes, even to believe what they are told. But this has nothing to do with the conventions of language.

## 6. *Crisscrossing Conventions*

Another thing that the rules describing conventional patterns do not mandate is that these patterns or portions of them cannot be used outside the conventions. Broken fragments of conventional patterns are frequently used purposefully in secondary ways. They are literally or figuratively 'played with' in ways requiring only that the participants understand from what wholes they were taken. Compare uses such as in word games, irony, metaphor, and so forth. Pieces having the same form as those used in one conventional pattern may be used also in constructions that have nothing to do with that convention. Sometimes they are portions of other unrelated conventions. Raising one's hand is conventional in order to request to speak, but it is also conventional in order to vote. That it is conventional to use the shape or sound 'bank' to refer to river banks does not prohibit its use in another convention to refer to financial institutions, nor does it prohibit its use in nonsense rhyme.

Where conventions overlap, using some of the same forms, it is not a matter of convention which if any is being instanced on a given occasion. That a token of any form instances a convention or piece of a convention is a matter of its individual history, not a matter of what it matches. The rules of chess do not tell when moving my queen next to your king constitutes playing that she kisses your king instead of putting your king in check, nor do any other conventional rules do this job. Only tokens reproduced due to weight of precedent are conventional, and which convention each instances depends on the precedent from which it is derived. Which if either of the hand-raising conventions is instanced when a particular hand goes up depends on which if any previous instances of hand-raisings are the causes and models for this one—whether instances used for requesting to speak or instances used for voting. Which, if either, conventional word 'bank' is being used depends on which family of previous instances this instance was copied from, not on some meta-convention or rule.<sup>10</sup>

<sup>10</sup> Compare my (1984: ch. 4) and also Kaplan (1990).

A language consists in a tangled jungle of overlapping, crisscrossing traditional patterns, reproducing themselves whole or in part for a variety of reasons, and not uncommonly getting in each other's way. Places where these patterns cross can produce ambiguities. These are sorted out not by conventions, but by the hearer managing to identify, by one means or another, the source of the pattern; that is, from which family it was reproduced. Nothing guarantees recognition of these sources. There cannot be, for example, such things as absolute semantic markers in a language; that is, forms reserved *merely by convention* to serve always the same function. No convention can prevent the possible accidental incursion of other crossover conventions. Where conventions cross, that the conventional pattern piece actually instanced is usually identifiable by the hearer is a function, mainly, of the speaker's skill, not of any conventions or rules of language. It is not conventional, for example, but surely good sense, not to raise your hand to speak when a hand vote is being counted. True, there can be conventions that ride piggyback and concern the use of other conventions. To avoid people taking back their hand votes later by pleading lack of intent, a rule might be made or another convention arise that when a hand vote is taken, no matter why a person put up his hand, it is counted. Still, as Dan Sperber and Deirdre Wilson (1986) have rightly insisted, understanding a language is never just decoding.

## 7. *Conventional Acts and the Speech-Act Tradition*

As I have described natural conventions, a conventional act is an act of reproducing or helping to reproduce, in whole or in part, a conventional pattern. Thus, people who reproduce conventional patterns do things that have been done before, and presumably were once done for the first time, thus done at least once without precedent. Conventions, I have claimed, really are nothing whatever beyond reproduced patterns of perfectly ordinary activity. Perfectly ordinary activities, like putting a letter in the mailbox and then putting up the flag, or taking a letter out

of the box after seeing a flag, are none of them such that they could not in principle, perhaps accidentally, occur in the absence of conventions. It follows that conventional acts achieve natural ends, but achieve them in conventional—that is, in reproduced—ways. I can set my table, hold my fork, and use my napkin in conventional ways or in unconventional ways. This was apparently the way H. P. Grice was thinking about acts of communication in 'Meaning' (1967/71). I was following a different but parallel route earlier when I described the possible evolution of fully conventional leader–follower coordinations from exaggerated and stereotyped natural anticipatory movements.<sup>11</sup> Given this sense of 'conventional act', proceeding with a bridge game in the conventional manner after someone has said 'I bid two diamonds', and also believing what one is told and doing what one is told (these latter being traditionally counted among the 'perlocutionary acts') are all conventional acts in exactly the same sense as are, say, conventional biddings and tellings, traditionally labeled 'illocutionary acts'.

The tradition of speech-act theory has something quite different in mind with the term 'conventional act'. John Searle (1983: 178–9), for example, tells us that, unlike illocutionary acts, perlocutionary acts are never conventional. From J. L. Austin (1962) through P. F. Strawson (1964), Searle (1969), and Stephen Schiffer (1972), to Kent Bach and Robert Harnish (1979) and Donald Davidson (1984), this tradition has taken it that 'conventional acts' are acts the performance of which would be strictly impossible in the absence of conventional rules. Nor is it meant, merely, that you cannot do a thing in a conventional way unless there exists a conventional way to do it. Rather, there are supposed to be conventions that 'constitute' certain natural acts (for example, raising your hand) as being, also, different acts that are conventional acts, such as acts of voting or requesting to speak, or of giving someone orders, or of asserting something, or of christening a ship, or of making a motion in a meeting, or of marrying someone. Thus, Bach and Harnish (1979: 109), closely following Searle (1969: 51–2), claim that a convention requires that there be circumstances such that it is mutually believed in some group that whenever such-and-such action is performed in

<sup>11</sup> A critique of Grice's way of doing this is in my (1984: ch. 3).

those circumstances, a certain conventional outcome is the result: for example, the action is automatically 'counted as' a marrying, or a christening. Strawson says of conventional speech acts, 'the speaker's utterance is not only *intended* to further, or affect the course of, the practice in question in a certain conventional way; in the absence of any breach of the conventional conditions for furthering the procedure in this way, it cannot fail to do so' (1964: 612). On these grounds, Strawson and Bach and Harnish deny of many illocutionary acts that they are conventional, replacing convention with Gricean intentions in the analysis of these speech acts. Davidson actually denies that any language is conventional on the grounds that its (ordinary) uses do not result in conventional acts of this sort (1984: 276). I close this chapter by offering a description of phenomena generally taken to be paradigms of 'conventional acts' in this latter strong sense, but a description that shows them instead to be constituted merely by natural conventions of the kind I have been describing.

These seemingly extraordinary acts (voting, christening someone, giving orders), I shall argue, reflect little more than a way that we have of *classifying* some kinds of conventional, and also some kinds of regulated, patterns of social behavior. Examine any part forming the beginning, or occurring in the middle, of some conventional pattern of activity. There will be a piece or pieces of the pattern that conventionally follow after this part. Or perhaps this part merely puts constraints on what can follow while conforming to the convention. Call the chosen part a *conventional move*, and call what must follow its *conventional outcome*. A move while playing chess or solitaire is a conventional move in this sense. It places restrictions on what can, conventionally, happen after. Telling someone to do something is a conventional move. Its conventional outcome is that that person complies. Saying 'the meeting is adjourned' in the right circumstances is a conventional move. Its conventional outcome involves, for example, that no more motions are made and no more votes are taken.

Besides conventional patterns of activity, there are other patterns that are regulated by law or other authority: for example, the patterns of activity that constitute applying for citizenship or for a visa and its being granted or not granted, and the patterns involved in entering into

a legal contract along with what legally follows. Call any parts of such patterns that have mandated outcomes or outcomes that are constrained by authority *regulated moves*, and call the mandated or constrained outcomes *regulated outcomes*. There are also many moves that are hybrid between being conventional and being regulated, a good example being traditional marriage ceremonies. Certain kinds of conventional-move outcomes, and likewise many hybrid and regulated outcomes, consist in the honoring by the society at large or its representatives of certain rights and privileges, or the carrying out of certain acknowledged duties; that is, more generally, in the imposition of certain restrictions on the behaviors of people designated as involved.

Now, often there are numerous different existent conventional or regulated patterns, reproduced or produced in different cultures or perhaps in the same culture, which contain differently constituted initial or middle portions, but which have the same or similar conventional or regulated outcomes. Then we are likely to possess verbal tools for classifying these conventional or regulated moves according to outcome rather than structure. Rather than talk of the natural actions that are initial or middle parts of these conventional or regulated patterns, such as raising one's hand or uttering certain sounds, we speak of voting (how?), or saying that it is raining (in what language?), or ordering someone to leave, or performing a marriage ceremony, or paying back a debt, or appointing a committee, and so forth.

In such cases, because the actions are described or identified only by reference to their conventional outcomes, of course they must be actions that are 'constituted as what they are'—that is, correctly placed under these labels—due to the existence of certain conventions or regulations. They are, as such, *essentially conventional* (or *essentially regulated*). Still, there is nothing more here than perfectly ordinary natural convention and perfectly ordinary social regulation—nothing, that is, except terminology. These acts are still just ordinary actions that have been reproduced after a pattern, or following a regulation, in no way any more or differently conventional than is the act of shaking hands with the right hand, or believing it when you are told it is raining. Throwing this sort of terminology into the passive, we say that people have been appointed to committees, or are married, that the meeting is



now adjoined, and so forth. That is, conventional patterns have been initiated that have certain sorts of conventional outcomes. (That these patterns have these conventional outcomes does not imply, of course, that they always or even usually come out that way.)

These oversimplified remarks may make some start, though rather a small one, toward understanding what speech acts are, and what their relations are to language conventions. Later (Ch. 8; Ch. 9) I will argue that the central criterion that we use in categorizing speech acts does not, in the first instance, concern the conventional outcome of the act but rather the purpose of the words uttered. This purpose often coincides, of course, with the conventional outcome. The purpose of a spoken set of words, however, has two independent sources: (1) the intention of the speaker, and (2) the purpose or proper function, in the public language, of the expressions used. In central cases both of these purposes will coincide with the conventional outcome of use of these expressions, but any two or all three of these can also come apart.

## REFERENCES

- Austin, J. L. (1962), *How to Do Things with Words* (New York: Oxford University Press).
- Bach, Kent, and Harnish, R. M. (1979), *Linguistic Communication and Speech Acts* (Cambridge, Mass.: MIT Press).
- Davidson, Donald (1984), 'Communication and Convention: Moods and Performances', in *Inquiries into Truth and Interpretation* (New York: Oxford University Press), 265–80.
- Gilbert, Margaret (1983), 'Agreements, Conventions, and Language', *Synthese*, 54: 375–407.
- (1989/92), *On Social Facts* (New York: Routledge; repr. Princeton, NJ: Princeton University Press).
- Grice, H. P. (1967/71), 'Meaning', *Philosophical Review*, 66/3: 377–88; repr. in J. Rosenberg and C. Travis (eds.), *Readings in the Philosophy of Language* (Englewood Cliffs, NJ: Prentice-Hall), 436–43.
- Kaplan, David (1990), 'Words', *Aristotelian Society Proceedings*, suppl. vol. 64: 93–119.
- Lewis, David (1969), *Convention: A Philosophical Study* (Cambridge, Mass.: Harvard University Press).

- 
- (1975/96), 'Languages and Language', in K. Gunderson (ed.), *Language, Mind and Knowledge* (Minneapolis, Minn.: Minnesota University Press); repr. in H. Geirsson and M. Losonsky (eds.), *Readings in Language and Mind* (Oxford: Blackwell), 134–55.
- Millikan, R. G. (1984), *Language, Thought, and Other Biological Categories* (Cambridge, Mass.: MIT Press).
- Recanati, François (1987), *Meaning and Force: The Pragmatics of Performative Utterances* (New York: Cambridge University Press).
- Schiffer, Stephen (1972), *Meaning* (New York: Oxford University Press).
- Searle, John (1969), *Speech Acts: An Essay in the Philosophy of Language* (New York: Cambridge University Press).
- (1983) *Intentionality: An Essay in the Philosophy of Mind* (New York: Cambridge University Press).
- Sperber, Dan, and Wilson, Deirdre (1986), *Relevance: Communication and Cognition* (Cambridge, Mass.: Harvard University Press).
- Strawson, P. F. (1964/71), 'Intention and Convention in Speech Acts', *Philosophical Review*, 73/4: 439–60; repr. in J. Rosenberg and C. Travis (eds.), *Readings in the Philosophy of Language* (Englewood Cliffs, NJ: Prentice-Hall), 599–614.
- Tinbergen, Niko (1952), '“Derived” Activities, their Causation, Biological Significance, and Emancipation during Evolution', *Quarterly Review of Biology*, 27: 1–32.

## 2 In Defense of Public Language\*

---

a notion of 'common, public language' that remains mysterious . . . useless for any form of theoretical explanation . . . There is simply no way of making sense of this prong of the externalist theory of meaning and language, as far as I can see, or of any of the work in theory of meaning and philosophy of language that relies on such notions, a statement that is intended to cut rather a large swath.

(Chomsky 1995: 48–9)

It is a striking fact that despite the constant reliance on some notion of 'community language' or 'abstract language,' there is virtually no attempt to explain what it might be.

(Chomsky 1993: 39)

either we must deprive the notion communication of all significance, or else we must reject the view that the purpose of language is communication . . . It is difficult to say what 'the purpose' of language is, except, perhaps, the expression of thought, a rather empty formulation. The functions of language are various.

(Chomsky 1980: 230)

I have yet to see a formulation that makes any sense of the position that 'the essence of language is communication.'

(Chomsky 1980: 80 (see also 1992b: 215))

At frequent intervals over the years Professor Chomsky has inveighed against both commonsense and technical notions of public language or

\* Reprinted with permission from N. Antony and L. Hornstein (eds.), *Chomsky and his Critics* (Oxford: Blackwell, 2003), 215–37.

'externalized language', claiming that they are confused, ill-defined, or of no scientific interest. As a scientist, he would be interested in public language only if it were a 'real object of the real world' (Chomsky 1993: 39) rather than an 'artifactual' and 'arbitrary' notion (Chomsky 1985: 26). I propose to articulate such a notion of public language for him. Chomsky has also denounced the notion that the purpose of language is communication. I will argue, on the contrary, that a primary function of the human language faculty is to support linguistic conventions, and that these have an essentially communicative function.

Despite this bold display of sounded disagreements with Chomsky, however, I agree with his objections to common views of public language completely: 'People who live near the Dutch border can communicate quite well with people living on the German side, but they speak different languages in accordance with the sense of the term Dummett argues is "fundamental"' (Chomsky 1992a: 101). Public language is a sprawling mass of crisscrossing, overlapping conventions, some known to some people, others to others. It does not divide into discrete portions, German versus Dutch versus French.

What we say is that the child or foreigner has a 'partial knowledge of English,' or is 'on his or her way' toward acquiring knowledge of English, and that if they reach the goal, they will then know English. Whether or not a coherent account can be given of this aspect of the common sense terminology, it does not seem to be one that has any role in an eventual science of language. (Chomsky 1985: 16)

If Peter is improving his Italian or Gianni is learning his . . . [w]e gain no insight into what they are doing by supposing there is a fixed entity that they are approaching, even if some sense can be made of this mysterious notion. (Chomsky 1992b: 16–17)

Languages are not governed by discrete sets of public norms or rules of mysterious origin. There are no 'fixed entities' which a child slowly approaches when learning language. On both the above points I entirely agree with Chomsky and entirely disagree, for example, with Dummett.

This 'externalized language' that Jones and Smith share must be an abstract object of some sort, a property of the community, perhaps . . . Suppose that Smith and Jones have more or less the same shape; we do not conclude that there

is a shape that they partially share, and the interactions between Smith and Jones give us no more grounds to suppose that there is a language that they share. (Chomsky 1993: 39–40)

Further, public language is not just some property shared by Smith and Jones. It is not discovered, for instance, by averaging over the idiolects of people in ‘the community’. I agree with all these points against the existence of what some conceive as ‘public language’. I also agree with Chomsky’s arguments for the importance of the study of I-language and the language faculty. And of course I agree that language is put to many different purposes. What, then, am I complaining about?

I will argue that there remains a legitimate way of looking at language as a public object, and also a legitimate way of looking at language function, that Chomsky has not taken into account. Learning language is not merely acquiring an ‘I-language’. It is not just achieving a relatively steady state of the language faculty. Learning language is essentially coming to know various *public* conventions and, with trivial exceptions, these conventions are around to learn only because they have functions.

What has brought Chomsky to deny these rather mundane truths, hence to embrace an unnecessarily extreme position on public language, may have been, in part, bad philosophical company. In particular, there are, I believe, two dominant but mistaken traditions in philosophy that probably played a role. One of these concerns the nature of language conventions. The other concerns the nature of language functions.

A ‘convention’ in Lewis’s sense, is a regularity ‘in action or in action and belief’ sustained by the belief that others conform to the regularity. Note that this is a rather restricted sense of the term ‘convention.’ There are, no doubt, conventional aspects of language: for example, the fact that one says ‘Hello’ in answering a telephone or calls a chair ‘a chair.’ But . . . regularities in action and belief are quite restricted, at least if we insist that ‘regularities’ have detectable probabilities; there is little reason to suppose that aspects of language that are commonly called ‘conventional’ involve detectable regularities. (Chomsky 1980: 81, referring to Lewis 1969)

Chomsky is completely right about this. The conventions of language are not regularities, either *de facto* or *de jure*. They are not expressed in actual uniform behaviors, nor in people’s beliefs about uniform behaviors,

nor in rules or norms prescribing uniform behaviors. But it is exactly because conventions have been taken, mistakenly, to be regularities (Searle 1969; Lewis 1969, 1975; Schiffer 1972; Bach and Harnish 1979; Gilbert 1983, 1989/92; Recanati 1987)<sup>1</sup>—and because language is so obviously conventional—that it has been thought that languages must belong to circumscribed social groups. You can't have a regularity unless there is a reference class for the regularity to be in. So there can't be a convention unless it is a convention for somebody—for the Jews but perhaps not for the Christians, for the Germans but perhaps not the French, for you but perhaps not for me. Hence the notion of a "common, public language" that remains [so] mysterious', of a "community language" or "abstract language" ', a 'shared' language, 'a property of the community'. Chomsky sees the error in this way of thinking. But the philosophical tradition has supplied no alternative to this way of understanding conventions, hence no alternative reason to think the study of public conventions might be important in the study of language. I will try to remedy this situation by outlining a theory of language conventions that does not take them to be regularities, and then showing why this sort of public conventionality is of the very essence of human language. I know nothing about the language faculties of angels or Martians. They may come with all necessary language wired in so that they do not need to have language conventions. But in humans, I will argue, a pivotal job of the language faculty is to make language conventions possible, and the functions of language conventions are communicative functions.

The second philosophical tradition to which I will propose an alternative concerns the nature of language functions:

What does it mean to say that language has [communication as] an 'essential purpose'? Suppose that in the quiet of my study I think about a problem, using language, and even write down what I think. Suppose that someone speaks honestly, merely out of a sense of integrity, fully aware that his audience will refuse to comprehend or even consider what he is saying. Consider

<sup>1</sup> Recanati (1987) takes it that conventional language devices 'indicate' or, using the linguists' term, 'mark' uses of language; that is, conventions mandate that these devices shall be used only for those purposes. This is not explicitly stated, but see, for example, sect. 22.

informal conversation conducted for the sole purpose of maintaining casual friendly relations, with no particular concern as to its content. Are these examples of 'communication'? If so, what do we mean by 'communication' in the absence of an audience, or with an audience assumed to be completely unresponsive, or with no intention to convey information or modify belief or attitude? (Chomsky 1980: 130)

Chomsky follows the dominant philosophical tradition in assuming that language functions would have to reduce to or be derived from speaker intentions. In fairness, the only alternative that mainstream philosophy has offered is to derive them somehow from language conventions, language conventions being understood, as above, to involve some obscure kind of regularity within some nebulous group. Chomsky is right to reject this latter alternative, but there is no need to accept the former. I will offer a third alternative, that I hope he will find more palatable. But given this third alternative, it will become clear that the function of the overwhelming majority of conventional language forms is to enable various kinds of communication.

The argument will be, then, that a central function of the language faculty in humans is to make language conventions possible, and that the functions of conventions are to make communication possible. In this case, I take it, 'has as its function to' is a transitive relation. At the same time, I will argue that the web of conventions that forms the mass that is public language is not an abstract object but a concrete set of speaker–hearer interactions forming lineages roughly in the biological sense. These lineages and their interactions with one another are worthy of scientific study. Nor are their properties derivative merely from the properties of I-languages.

Both the thesis on conventions and the thesis on functions to be applied here have been fully stated and defended in other places.<sup>2</sup> Here I can only sketch results and try to ward off the most obvious objections. First, then, some words about conventions.

As Chomsky remarks, it is conventional to say 'Hello' when you answer the phone. Moreover, nearly everyone, say, in America, does so.

<sup>2</sup> The thesis on conventions is in Ch. 1; the thesis on language functions is in Millikan (1984, 2001, ch. 8, ch. 9).

But it is conventional in exactly the same way to call a spigot 'a spigot' and also to call it 'a faucet' and also 'a tap', and clearly there can be no group in which everyone does all three of these things regularly. And of course that is not what David Lewis had in mind when he claimed that language conventions exist only where there are regular behaviors. He didn't mean that spigots are always called 'spigots'. What he meant was that if someone calls something 'a spigot', then it is regular that that thing is a spigot, which looks quite a lot more plausible. Troubles with regularity theories show up not here but, first, with the problem of how to delimit the relevant groups within which the regularities supposedly occur. Within what designatable group is the convention to hand out cigars on the birth of a boy regularly followed other than the group that in fact hands out cigars on the birth of a boy? Within what group is the convention that by saying 'break a leg' one wishes another luck regularly followed other than the group that in fact uses 'break a leg' to wish people luck? Troubles show up with the regularity theory, second, where people use language in ways that do not accord with convention; for example, when they use metaphors, sarcasm, and other figures of speech, and when they lie, or make false promises, or demand things they know they won't get, or, as Chomsky suggested (though in another context), when they sincerely assert what they know won't be believed, or when they talk only to themselves. There is a voluminous literature on such cases in speech-act theory, which tries vainly to solve an insoluble problem. It tries to describe regularities involved over all such uses—very subtle regularities to be sure, usually taken to occur on the level of multiply embedded beliefs that speakers supposedly invariably intend to impart. Since Chomsky apparently agrees that this sort of problem is insoluble, I won't argue the case here. I will simply start in and describe a kind of convention that has no connection with regularities, that is not a convention 'for' any particular group. And I will try to show that this is a way in which natural human language is conventional.<sup>3</sup>

What is conventional in this sense is an activity or a pattern of activities. Exhibition of the pattern may require only one participant or it may involve more. Conventional ballroom-dancing steps, for example,

<sup>3</sup> For more detail see Ch. 1.



are patterns involving two participants, whereas the pattern that is playing ring-a-ring o' roses generally requires more. To become conventional, an activity or pattern of activity must, first, be reproduced, hence proliferated. (Please do not blanch. 'Reproduced' will be used here such that language forms can be reproduced by a language faculty imposing a universal grammar.) Further, it must be proliferated due in part to the weight of tradition, rather than due, for example, to its intrinsically superior capacity to perform some function, or due to ignorance of any superior way to perform it. I will discuss each of these characteristics of conventions in turn.

First let me explain 'reproduced'. I will use this term in a limited way. A reproduction must be such that had the original been different in specifiable respects the reproduction would have differed accordingly. Inherited traits and behaviors are not 'reproduced' in this sense. I have blue eyes not directly because my mother and/or father had, but because of my genes, which were copied from their genes, which were not, however, copied from their eyes. Had my mother's or father's eyes been yellowed from jaundice, that would not have caused my eyes to be yellow.

A reproduction is always a reproduction only in certain respects. The reproduction that comes out of a monochrome copying machine is a reproduction only with respect to pattern of light and dark. The background color and the paper texture are not reproduced, as these depend on the color and texture of the paper put in the paper feeder, not on the color and texture of the original. The color of the pattern is not reproduced because it depends on the color of the ink in the machine, not on the color of the original. Any object can be reproduced in an indefinite variety of ways, since any object has an indefinite number of properties. Nothing is a reproduction of anything else in all respects. Because properties that are reproduced can be relational properties as well as intrinsic properties, and since everything has an infinite number of relational properties, it is not possible to exhaust the variety of reproductions that might, logically, be made of an object. Reproductions that copy relations may be unlike their originals in very striking respects. For example, a painter who reproduces 'the style' of an earlier painter may paint quite different subject matters and use quite different media,

while another painter may reproduce this same painter's 'style' by painting similar subject matters only. The products of both painters will be 'reproductions' as I am using that term. Further, since an object can have many parts and many aspects or properties, it can be a reproduction of a number of different things all at once, borrowing some features from here, others from there, and so forth.

For one thing to be 'reproduced' from another, all that is required is that there be a mechanism that produces the second on the model of the first, such that had the first been different in specifiable respects, that would have caused the second to differ accordingly. Under this description, reproducing can of course occur unconsciously. For example, there can be conventional ways of talking or of moving the body, differing for men and women, or specific distances from one another at which it is conventional to stand when talking, that are handed down by being unconsciously reproduced. Also, reproducing need not be direct. If Mother tells Johnny to shake hands when being introduced, rather than, say, telling him to kiss on both cheeks or to sniff noses, granted that she tells him this only because other folk behave or have behaved in this way, then Johnny's behavior, when he does what he has been told, is indirectly reproduced. Had certain other people been greeting one another in other ways, Johnny's way of greeting would have differed accordingly as a result. Also, if Jim taps a nut to fit Jon's bolt, then Jan makes a new bolt to fit Jim's nut, on this definition, the thread on Jan's bolt has been reproduced from Jon's bolt. I will call this last kind of reproduction *nuts and bolts reproduction*. Nuts and bolts reproduction probably figures in the propagation of certain kinds of conventional language patterns.

Now a preliminary word about the relation of function to convention. Many conventions have no apparent function. It is conventional for a bride to wear a blue garter, and to throw her bouquet to the bridesmaids. It is conventional to dress girls in pink and boys in blue, and to put a wreath on the door at Christmas. But some conventional patterns obviously do serve functions. For example, the convention of driving on the right in the USA serves a function. Where a reproduced pattern does serve a function, it is not a conventional pattern unless it is one that would have no particular reason to emerge again, rather than some

alternative pattern, if once forgotten. This is because its intrinsic nature makes it no more able than other known or equally knowable patterns to serve its function. Conventional patterns are exemplified rather than other patterns owing only to historical accident, but having occurred, they cause their own recurrence.

A convention of this kind is not tied by definition to any particular group, nor is there anyone in particular who must know of it or follow it. Obviously, there have to be some people who follow or have followed it or the pattern is not reproduced and hence not conventional, but it is not necessary that any particular people should do so, nor that a large number of people should do so. When conventions are associated with specific groups, this is because it happens to be the people in these groups who know the conventions and happen, for whatever reasons, to follow them. Which independently designated 'group' knows that it is conventional to drink green beer on St Patrick's Day? How many who know accord? What else do these people have in common that would form them into a social group or 'community'?

A convention does not prescribe that everybody in some predesignated group should follow it. To know a convention is not to know what to do categorically, but only to know what to do if you wish to follow the convention. For example, the conventions, the rules, of chess do not tell you what to do, but only what to do if you wish to play chess. You can get these conventions wrong, of course. You can fail to reproduce the conventional chess patterns faithfully even though you intend to. But the standard that has then been violated was set by your own intention. No public prescriptive rule will have been violated.

The appearance that conventions involve prescriptive rules may result partly from the fact that many conventions are conditional. The convention is to do something in a certain kind of context, or at a certain time, or if you are a certain kind of person, or in a certain situation. The convention is to put up red and green decorations at Christmas time, or to say 'Hello' when you pick up the phone, or to wear a ring on a certain finger of your left hand if you are engaged or married. But that the conventional pattern includes a conventional setting does not mean that it mandates itself. A convention is not, as such, a rule that is required to be followed. A convention is not something having magic

times or places or situations where it is mysteriously 'in force'. If there is some kind of mandate or compulsion to follow a convention (legal, moral, a norm of etiquette) or to follow it under certain conditions, that is a separate thing entirely, and not what makes it a convention, at least not in the sense of 'convention' that applies to natural language forms and patterns.

Nor—and this will later emerge as crucial to the case of language conventions—does a convention mandate that its pattern must be finished once begun. More generally, it does not mandate that the whole of the pattern must be reproduced and not its parts separately. One can quit a chess game in the middle, or set up just an end game or a middle game. Although there is a convention to have one's soup first, then one's salad, main course, and last one's dessert, one can follow the convention part way but not have dessert. For all this kind of convention is, basically, is a pattern of activity that gets reproduced.

Soon I will discuss conventions that serve to coordinate actions and thoughts between people, for these are very important to an understanding of language. And I will argue, *contra* Lewis (1969), that like more simple conventions, coordinating conventions too do not, in general, involve regularities. First, however, we should examine some simpler aspects of public language in the light of its conventionality. We can begin with words.

Words, many philosophers have thought, are typed or individuated in accordance with physical form. My guess is that there are no legitimate theoretical purposes for which this way of typing is relevant, but I will not argue that here. Rather, I will examine some ways of typing tokens into 'words' on which we commonly rely in ordinary contexts, and which reveal our ordinary sensitivities to the conventional aspects of language as I have described them. According to one way of typing, there are parts of Tennessee, for example, where the word 'pen' sounds exactly the way my word 'pin' sounds. It's exactly the same word, but it sounds different in Tennessee than Connecticut. 'They say many of their words differently in Tennessee.' Clearly the same word again need not be made of the same acoustic sounds again. Linguists may say that it must be made of the same phonological segments again, the same phonological segment sometimes taking a variety of forms, there being,

for example, different allophones of it. On a similar way of typing, the same word must be made of the same letters again, but the very same letter has quite different shapes when different people write it, while different letters sometimes have the same shape, for example Greek P (rho) and Roman P. Typed in these ways, word types and their elements, phonological segments and letters, are like species.<sup>4</sup> In biology, what makes a dog a dog is, in the first instance, that it was born of a dog, not that it has some particular shape. Similarly, what makes a shape or sound into a token of a particular word on this way of reckoning is its lineage, what it was reproduced from, on what prior word tokens it was modeled, or alternatively (and this can be different), on what combination of phonological segments or letters, these being typed by their histories, their lineages. The lisping child who pronounces the word 'red' like 'wed' does not pronounce the word 'wed' but mispronounces the word 'red'. If I repeat after a Scotsman, I say the same words, but many of them sound quite different. The same word is spelled 'color' in America and 'colour' in England. Conversely, the twinearth word 'water' is as much a different word than ours as the stuff to which it applies is a different liquid than ours. It is different the way Greek P (rho) is different from Roman P. And, of course, comments about the river bank and about the bank and trust contain different words 'bank', copied from completely different lineages. These ways of typing word tokens depend on the conventional nature of words, on their reproductive nature, in much the same way that the typing of individual animals depends on the reproductive nature of animals.<sup>5</sup> And they run into similar problems. How much can a species change over time without becoming a different species? Is it really so that the word 'moan' in 'He moaned and groaned' is the same word as 'mean' in 'He didn't mean to' but a different word than 'mean' in 'no mean city', as the etymologist is liable to say?

<sup>4</sup> I realize that this is not how contemporary linguists *define* phonological segments, but I am suggesting this as a theory concerning an aspect of their real nature. Similarly, biologists before Darwin thought that all members of the same species had an inner nature or form in common. Modern biologists don't think so.

<sup>5</sup> This point about the typing of words is put in a broader context in Millikan (1984: ch. 4). Much the same point about words is also argued in Kaplan (1990).

Similarly, the difference between two tokens of a word being used in two different senses and being used once in a literal and once in an extended way is a question of lineage. Word tokens used in extended ways are reproduced from words used literally. Tokens of the same word that have taken on different senses are words with a common lineage some distance back, but whose lineages have now separated—a matter which can, of course, be one of degree. For example, the question of whether the word ‘red’ in the phrase ‘red hair’ has a different sense than the ‘red’ in ‘red dress’ depends on the ancestry of its tokens. Are its tokens currently being reproduced only from other tokens of ‘red hair’? Or do people individually and independently each stretch the word ‘red’ as used in contexts like ‘red dress’ to cover also red hair? Considered in this light, the conventionality, hence public nature, of words is certainly a real and interesting phenomenon, though often hidden from direct observation in its details.<sup>6</sup>

Besides words, grammatical forms are conventional patterns in the sense we are explicating. Remembering that a reproduction always copies only certain aspects of its model, clearly children reproduce syntax on the model of what they hear. Nor do I mean, of course, that grammar is a matter of patterns in sound. It is the linguist’s job to discover exactly what kinds of patterns are materials for the conventions of grammar, to discover what forms of what aspects of speech are the reproduced ones in conventional grammars. The hypothesis of a universal grammar faculty that imposes restrictions on human grammars concerns this matter. Whatever the details, a UG, looked at one way, acts as a more or less complicated filter governing which aspects of what is heard will be reproduced, which aspects will vary depending on models heard and which will not. If there are certain aspects of grammar that are never reproduced at all, but always resupplied by the language module *de novo*, these act like the color of the ink cartridge in the monochrome copying machine. They are not conventional aspects of grammars. A huge question for linguists, of course, concerns which aspects of grammars are and which are not conventional, in exactly this sense of ‘conventional’.

<sup>6</sup> There is more on all the above matters in Millikan (1984: chs. 3, 4) and in Ch. 3 below.

It is important to notice the similarity here between grammatical conventions and others. The perpetuation of any convention requires, categorically, that participants attend to the same aspects of the convention's embodiments and not to others, so that they generalize to new cases in the same manner. Thus, the convention is to hang a wreath on the door at Christmas, not a sweet-smelling object, and not to do it at relative-visiting times, or at present-giving times, or on cold Tuesdays. In the case of every convention, there has to be some mechanism, whether accidental, or whether designed by nature or man, that operates to keep copies and copies of copies somewhat uniform. Otherwise there is, of course, no convention at all.

A UG is a mechanism effecting uniform reproduction of syntactic forms, hence the maintenance of syntactic conventions. Maybe UG got there accidentally, as Chomsky has sometimes suggested, or maybe it got there by natural selection; it doesn't matter. My present point is that without something like it there could be no grammatical conventions. Similarly, if there are inborn mechanisms, as there appear to be, that efficiently accomplish mastery of the phonological structures of languages that the young child hears, these function as another kind of filter that serves to narrowly channel linguistic reproduction, hence to aid the proliferation of linguistic conventions. Phonological mastery of a language yields a generalized sameness-difference schema for the language, dictating what is to count as another correctly executed linguistic utterance of the same type, along one relevant dimension of linguistic reproduction. Alvin Liberman has argued that phonological structure is the fundamental *sine qua non* allowing for the practical possibility of language innovation (Liberman 1999). Without it we could at best be stuck with a limited inborn vocabulary, that had slowly and painfully accrued during the course of genetic evolution. A great deal of attention has been paid to the kind of productivity made possible by a grammar that allows embeddings, so that an infinite 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, is productivity with a significance at least as profound.

This idea [that new speech forms that a speaker has not heard are produced on analogy with those he has heard] is not wrong but rather is vacuous until the concept of analogy is spelled out in a way that explains why certain 'analogies' are somehow valid whereas others are not . . . We can give substance to the proposal by explaining 'analogy' in terms of I-language, a system of rules and principles that assigns representations of form and meaning to linguistic expressions . . . but . . . with this necessary revision in the proposal, it becomes clear that 'analogy' is simply an inappropriate concept in the first place. (Chomsky 1985: 32)

'Analogy', in this context, could mean just sameness in abstract or relational form. Producing new speech forms on analogy with old could mean just some sort of reproduction, as I am using the latter term. But what people more usually have in mind when they speak of 'analogy' in this context is a fairly free sort of reproduction, that might pick up one aspect of the original or might pick up another. That language learning and use involve analogy in this free sense is something that both Chomsky and I want firmly to deny.

On the other hand, the present perspective on public convention gives us little reason to suppose that 'I-language' should consist in something appropriately described as 'a system of rules and principles', at least not if 'system' implies much systematicity. To view the young child's language faculty as a filter through which language conventions are to be transmitted is to view it not as aiming toward some steady state as the child matures, to view it not, for example, as a process of acquiring permanent parameter settings, but as a faculty engaged in the accumulation of a larger and larger repertoire of conventional patterns it can recognize and reproduce on demand. If it reaches a steady state, that will be only if it runs out of local conventions to learn. It would be likely to reach a steady state only if there really were such a thing as The German Language, The French Language, and so forth, to be learned. But there is quite a mass of conventions out there to be learned. These conventions are complex, and not particularly systematic, and



sometimes crisscrossing or contradictory, getting in one another's way. Thus, with Italian words you use Italian grammar, with English words English grammar, except that in some cases people do quite a bit of mixing, even in conventional ways. With 'expect' you use the infinitive but with 'anticipate' the gerund. With most English verbs you add an '-ed', but not with 'run' or 'swim'—people just aren't doing it that way here these days. And the conventions cross at 'Is the missionary ready to eat?' and (when it is spoken) 'Gladly, the cross-eyed bear'. Because linguistic conventions can be more or less compatible in various ways, and because they are often built on one another, exactly like genes in gene pools, they tend to get together in stable clusters. There are innumerable plant species that hybridize quite readily, but left to themselves in a relatively uniform environment the genes fall back again into stable clusters, and separate relatively homogenous species emerge again. In the same fashion, German and French and other languages are formed out of compatible strands of convention, and if different peoples are isolated from one another for very long, clear demarcations soon emerge between the language conventions they follow. But the idea that there is at the center of each such language some univocal 'system of rules and principles' seems as unlikely as that members of an animal species should be genotypically identical or that the whole gene pool should contain no incompatible genes. Possibly there is no real disagreement with Chomsky here. Maybe it is just that I am much more impressed—overwhelmed?—by the huge number of idiosyncrasies and idiomatic elements in any natural language. Perhaps whether 'a language' is best idealized as a monolithic structure, or as a relatively loose texture of interlocking crisscrossing conventions, is only a matter of what you are interested in. Surely either way idealization is involved. Chomsky and I agree, after all, that there actually is no such thing as 'a language'.

In any event, the phenomenon of public language emerges, I believe, not as a set of abstract objects, but as a real sort of stuff in the real world, neither abstract nor arbitrarily constructed by the theorist. It consists of actual utterances and scripts, forming crisscrossing lineages. What language forms one is using, from the standpoint of public language, depends not only on the settings in one's language module but on what public conventions one is following. Moreover, public language has

definite form, granted that it is passed on in a uniform way by people harboring I-languages that are definite in intension, as Chomsky requires.

If you take the gene's-eye view, Dawkins tells us (1976/89), the organism is just a gene's way of making another gene, but, of course, if you take the organism's viewpoint the gene is just an organism's way of making another organism. Similarly, if you are interested in individual psychology, public language is merely a stimulus to transition from the initial state of the language faculty  $S_0$  to a more steady state  $S_s$ . If you are interested in public language forms, on the other hand, the language faculty is merely how public language forms reproduce themselves. A public language is interesting in its own right, I will argue, because it has certain functions that are all its own, that are not merely abstractions gleaned by averaging over speakers' intentions. The study of the functions of public language is a separate discipline, independent of the study of individual psychology.

To show this will require making the second move that I mentioned at the start of this paper, introducing an appropriate theory of language functions. The functions of language conventions are, for the most part, coordinating functions. So I must also show how the existence of conventions, given the way I have described them, can produce coordinations. Coordinations effected through language conventions typically involve communication. The second main argument of this paper will then be complete: a primary function of the human language faculty is to support linguistic conventions, and these have an essentially communicative function. Therefore, a primary function of the human language faculty is to support communication.

I will be using 'function' here as short for 'direct proper function' as defined in Millikan (1984: ch. 1). This notion is a relative of the notion of function that biologists use when they distinguish between functions and mere effects of an organ or trait or activity. Roughly, the function or functions of a conventional pattern are those effects of it that account for its continued reproduction. More accurately, the pattern is proliferated due in part to a correlation between it and certain of its effects. It is selected for reproduction, in accordance with conscious or unconscious intent, owing to its being coincident with these effects enough of the time. Correlations can, of course, be either very high or very low. Many

biological items and traits have functions that they perform very seldom, yet just frequently enough to keep the genes responsible for them from drifting to extinction. Similarly, conventional activities and patterns of activity can have functions that they perform only once in a while, yet perform just often enough to keep them from becoming extinct.

Many conventions seem to have no functions. They seem to proliferate only because people are creatures of habit, or unthinking conformists, or because they venerate tradition, and so forth. Similarly, most patterns of activity that are reproduced due to their effects are not conventions, but rather are handed-down skills. Conventional *ways* of performing certain tasks, such as conventional ways of holding eating utensils in various cultures, or of sitting at meals, are proliferated, in a sense, because they serve certain functions. But it is only certain details of these activities that are conventional; namely, details not required to effect the functions that sustain the activities. Other kinds of details would do as well instead. On the other hand, some conventions have functions *as* conventions. The clearest and most interesting functions that conventions can have are coordination functions. I will call conventions with these functions *coordination conventions*.

Coordination conventions consist of patterns of activity (1) involving more than one participant, (2) proliferated because they serve a purpose had in common by the participants, (3) where the contribution to the joint pattern that each participant must make in order to reach the common goal depends crucially upon the contribution made by the other(s), and (4) where a variety of equally viable alternative joint patterns would achieve the same goal as well. Coordination conventions are thus patterns which might be said to be suitable to solving 'coordination problems', though the sense of the latter phrase would not be quite David Lewis's (Lewis 1969).

Some conventional coordination patterns require partners to do the same as one another, while others require them to do complementary things. Examples of patterns requiring the same include driving on the right, shaking hands with the right hand with an up-and-down motion, and standing at a conventional social distance when conversing. Examples requiring complementary actions include conventional

positions for partners of opposite sexes while dancing, the pattern *original caller calls back, original receiver waits* when a phone conversation is disconnected (Lewis 1969), the conventional pattern *driver on the right crosses, driver on the left waits* when arriving simultaneously at a four-way stop sign, and the pattern *patron puts the flag up when putting mail in, mailman takes the flag down when taking mail out* when a letter is mailed in a rural mailbox. The view that conventions require regularities has sometimes been supported in part by the assumption that all parties in a coordination pattern do the same thing. For example, Lewis describes patterns such as the telephone convention (above) as though both persons did the same, each following the same conditional rule, *if you are the original caller, call back, if not, wait*. But this is a vacuous move. Any pattern whatever involving more than one person can be described as though these people all did the same. Even the executioner and the executed do the same thing: *if you are the executioner, chop, if you are the executed, relinquish your head*.

For coordination conventions, it is important to consider whether in order actually to effect coordinations these conventions require regularities in the sense that all or most in some group should accord with them. The answer depends directly on how observable one partner's contribution to the pattern is at the time when the other or others must make theirs. Consider, first, the telephone convention. Here, each partner must make his or her own contribution while completely blind to what the other is doing. For this reason, the coordination will not be achieved unless each partner produces their part of the pattern before knowing whether the other will produce theirs. In this kind of case, use of the convention will be more or less effective depending on the frequency with which people who interact with one another conform to it, universal conformity producing the best results for everyone. Alternatively, each partner must have prior knowledge which of the various persons with whom they interact is likely to abide by the convention regularly, and if no one abides by it regularly, coordination cannot be achieved. Such a conventional pattern will produce coordinations often enough to encourage its own reproduction only if followed some critical proportion of the time.

The convention of driving on the right is almost as blind as the telephone convention, though not quite. One can see whether a car approaching ahead is driving on the right when the road is straight, but not, of course, around curves. Similarly, the mailman has no way of knowing whether the rural patron has performed his part of the conventional flag-raising pattern without looking in the mailbox to see if there is outgoing mail there. But the point of the convention is, of course, to avoid his having to look if the flag is not up, at the same time securing that all the outgoing mail gets collected. Again, these conventions will serve their functions better the more people follow them.

Contrast these cases of 'blind' conventions with the case of assuming conventional postures for social dancing followed by the joint execution of various conventional dance steps. There are a number of alternative conventional postures that a couple may assume, and many alternative combinations of conventional steps can be taken to a given piece of music. But here the conventional pattern initiated by one partner is immediately known to the other: the conventions are not blind but open. There is no need then for different dancers all to conform to the same postures and ordering of conventional steps. All that is needed is for one partner to lead and the other to follow. What will not work, of course, is for both to lead. And there is a convention that the man leads, but, again, not one that everyone has to follow. Conventionalized social dancing thus involves 'open' conventions rather than 'blind' ones.

Warming up to the case of coordinating language conventions, here are four more cases of open conventions: (1) I stand behind your car and wave my hands this way and that in a conventional way and you complete the conventional cooperative patterns by backing your car to suit my gestures; (2) While biking, I hold my left arm out and you stay out of my way as I make a left turn; (3) I bid four clubs and my partner bids his strongest suit (the 'Gerber convention'); (4) I want you to pass the bread whereupon I say, 'Please pass the bread', and you pass the bread, or, I want you to pass the bread whereupon I say, 'Brot, bitte', and you pass the bread.

With 'open' coordination conventions, there is no need for regular adherence to any one convention among others for achieving a given kind of coordination. Often many alternative open conventions coexist

quite compatibly. Recognizing that a leader has initiated a particular conventional pattern and completing it is like chiming in after the first line of a familiar song. Nor is it necessary that the partners who follow should be regular in their responses to conventional leads. Conventional patterns, as noted before, often persist even though quite regularly broken or interrupted in execution, partners being unwilling, or unable, or having other plans in mind, and so forth. If the coordination effected by completion of a conventional pattern even occasionally has enough value to the partners involved, the pattern may be able to survive even though more usually fractured than not. Perhaps the best contribution that Skinner made to psychology was his demonstrations of the effectiveness of random reinforcement schedules even when reinforcement is, on average, very infrequent. Lead portions of conventional patterns also are often turned to secondary purposes that do not accord with their functions *as* conventions; that is, with the reasons for survival of the conventional pattern. A dancer might lead into conventional dance patterns that he knows his partner can't follow in order to embarrass his partner or to show off. Or the one standing behind the car might jokingly signal a turn that both know leads into a brick wall.

I now want to make plausible that the functions of language conventions are primarily coordinating functions. It is primarily for the service of coordination between speakers and hearers that language patterns are selected to be proliferated as conventions. That is, were it not for the fact that employing its conventions sometimes serves purposes common to both speaker and hearer, language as we know it would shrivel and die. Indeed, for emphasis I might make a stronger claim. Putting completely to one side the evolution of the language faculty itself, were it not for their roles in the achievement of communicative coordinations, there is every reason to suppose that the individual language faculties of individual humans would atrophy, just as with unfortunate children who are not exposed to human language forms at all. Imagine, for example, a child who hears a perfectly normal assortment of English sentences, but for whom no coordinations are ever achieved through use of conventions involving these. Such a child, I am suggesting, would develop no language at all. That, of course, would be an empirical hypothesis, but I will offer reasons to entertain it seriously.

On the other hand, it is possible that there exist some language conventions that have no functions at all. Saying 'uhh' at intervals may be such a convention. The corresponding German convention is to say 'also', and I am told that in Hungarian one says 'öö'. Possibly the use of expletives has no coordinating function. The function of expletives, granted they have a function, may be simply to relieve oneself of, or to objectify for oneself, one's emotion. Thus, people use expletives just as often, perhaps more often, when alone as with other people. It is worth noting, however, that a different part of the brain is involved in the production of expletives than for other speech forms (Pinker 1994: 334).

The best argument that few if any language conventions lack coordinating functions is to discuss some that obviously do have coordinating functions. It will then become clear, I think, how ubiquitous such coordinating functions must be.

We can begin with functions of the grammatical moods. In English, at least, the syntactic forms identified with the indicative mood are proliferated in the service of a number of different coordinating conventions having different functions: conveying information or reminding someone of it ('The Athenians had slaves too'), giving orders ('You will report to the CO at 6 a.m. sharp'), conveying norms ('Johnny, we don't eat peas with our fingers'), and making declarations ('The meeting is adjourned') are four common ones, and undoubtedly there are more. The English indicative mood is like a homonym in this respect. Better, it is like a word that has several distinct though related meanings, instances of each use normally being modeled on prior instances of that same use, where all but one such use can still be heard as a dead metaphor. The use for conveying information is heard as the originating or most 'literal' use of the indicative, so let us look first at that pattern.

The pattern begins with a speaker S believing some proposition *p* and accordingly speaking an indicative-mood sentence that expresses *p*, given the truth-conditional part of the semantics of the language. It concludes when a hearer H, following the truth-conditional semantics for the language, translates the sentence into the thought that *p*, and accordingly believes that *p*. The pattern produces a coordination between speaker and hearer under the following assumptions: (1) the speaker is interested to convey information on the subject that *p* to the

hearer, (2) the hearer is interested to gain information on that subject, and (3) the proposition *p* is true. The end that *S* and *H* have in common is that *H* should become informed about *p*. For the entire pattern to have attainment of this end as its *own* or *proper* function, (1) it must be a reproduced pattern or (the normal case) it must be composed of reproduced patterns (the words, the syntactic form(s)), (2) it and/or its reproduced elements must each continue to be reproduced only because they sometimes make a certain definite contribution in the service of coordinations, and (3) the sum of these contributions, if made in this case also, would effect that *H* would become informed about *p*. For the indicative-mood form itself to have the transfer of information from speaker to hearer as a function, it must be that the contribution it makes to the whole is to utilize the functions of other reproduced parts of sentences exemplifying it to that end.

'Reproduction' in this case, and in the general case for coordinating functions of language, is mainly by nuts and bolts reproduction (see above, p. 31). The speaker parts of the patterns are reproduced due to the effects they sometimes have in the presence of hearers who complete them, and the hearer parts are reproduced due to the effects they sometimes have in the presence of speakers who have initiated them. But the whole pattern, involving speaker and hearer, is arbitrary in relation to its function; that is, it is conventional. *H* responds to the indicative sentence by translating it into belief in accordance with certain semantic rules because, in *H*'s experience, responding selectively to indicative sentences in this way has often enough resulted in the appropriation of useful information. *S* translates *S*'s belief into an indicative-mood sentence in accordance with these same semantic rules because, in *S*'s experience, often enough hearers respond to such sentences by forming beliefs accordingly. Read the 'because's' here not as indicating conscious reasons, but as indicating causes. That is, turning the coin over, had *H* not lived where speakers often enough expressed true beliefs using the indicative pattern with these truth-conditional semantic rules, *H* would not translate from indicative sentences into belief in this way, and had *S* not lived where hearers often enough translated from indicative sentences into belief in this way, *S* would not speak in this way when interested to convey information. Thus, speakers and



hearers collectively learn from each other how to speak and how to respond to speech in ways that serve purposes for both, each leaning on the settled dispositions of the other. The parallels with the evolution and fixation of symbiotic relations between animal species and with the evolution of animal signal systems should be apparent.

It should be clear from the above description that successful coordinations achieved in this way through use of the English indicative mood help not only to proliferate the indicative mood, but also to proliferate use of and reliance on semantic conventions applying to other elements of indicative-mood sentences. That understanding the truth-conditional semantics this way is enabling me often to collect useful information reinforces my use of these conventions for interpreting the semantics, and that my hearers often enough seem to believe what I intend them to believe reinforces my use of these semantic conventions when speaking. If I never had either of these experiences when using or interpreting indicative-mood sentences, and if, further, this use of the indicative mood was the sole convention in my language community involving intentional attitudes, hence the only one employing the truth-conditional part of its semantics, I could not possibly learn to understand either this function of the indicative mood or any of the truth-conditional semantic conventions of my language. Of course, the use of indicative forms to convey information is not the only convention employing the truth-conditional part of the semantics of any natural language. But the underlying principle here is crucial. It is only through various communicative uses of language and through the conventions that make these possible that the truth-conditional semantics of a language, which is also conventional, of course, is learnable. I will reinforce this point in a moment by discussing other kinds of conventions that rely on the truth-conditional semantic conventions of languages, and vice versa.

A surprise on this analysis of the conventional nature of the information-transferring function of the indicative is that believing what you hear said in the indicative turns out to be a conventional act, something one does in accordance with convention. Compare the following: standing, while conversing, at what happens to be the conventional social distance in one's culture is something one does in accordance with convention. And one learns to stand at the correct social distance in very

much the same way, by unconsciously learning to fit with others who are already standing at that distance. The difference is only that in the case of social distance the convention requires that partners do the same thing rather than different things. On the other hand, of course, the fact that H's believing *p* is a conventional outcome of S's telling H that *p* in no way mandates that H will or should believe *p*. The fact that a given outcome would be the conventional one in no way implies or mandates that it will be the actual one. As emphasized earlier, pieces of conventional patterns often occur independently. The speaker who lies also illustrates this point. He uses conventional structures in an attempt to induce partial completion of a conventional pattern, but both the beginning and the end of the pattern are missing. He does not translate from a genuine belief into words as required by convention, nor does the hearer end with a true belief as required.

Similar analyses can be applied to the functions of numerous other language forms, each of which continues to be reproduced due to its occurrence, often enough, in a pattern ending in a certain conventional hearer response. The root function of the imperative mood, for example, is to produce a corresponding action by the hearer, where the speaker is interested in having that action performed, and the hearer is interested in completing the conventional pattern, perhaps because he has a further interest in common with the speaker (hearers often want direction from speakers), or because conforming is sanctioned in one way or another. Performance of this imperative function is accomplished through the hearer's first forming an intention to perform the designated action, an intention formed following the truth-conditional semantic conventions of the language, so that reinforcement of the hearer's compliance tends to reinforce his observance of these semantic conventions as well. The imperative function also proliferates reproductions of indicative-mood forms, especially in the armed services. Here the indicative mood functions, as I have said, rather like a dead metaphor. The root function of the interrogative is to get information of a designated kind conveyed to the speaker, and its conventional use when successful also reinforces any truth-conditional semantic conventions used with it. And so forth. On the other hand, elsewhere I have discussed the functions of a number of other language devices (Millikan 1984), claiming that the functions

of sentences in which these forms occur do not directly implicate intentional attitudes on the part of either speaker or hearer. These sentences have as their functions to do other things, as it were, to hearer's heads. Similarly, Strawson (1972) claimed that what an identity sentence *does* is to merge two information files in the hearer's head (see also Millikan 1984, 1993, 1997), and Wilfrid Sellars (e.g. 1963: essays 4, 5, 6, 8, 10) claimed that what the 'X' means 'Y' rubric *does* is to prompt the hearer to use 'X' in the way he already knows to use 'Y'.

It is crucial that the functions of language forms are not the same as the uses to which they are usually put. Effective use of a set of language conventions requires the acquisition of a mass of skills and a good deal of inventiveness well beyond mere grasp of those conventions. Hearers can by no means be counted on to complete the conventional patterns that speakers initiate. They will do so only under special conditions. For example, they must trust the speaker, understand the subject matter, and have interests compatible with the coordination projected by the particular convention used. Moreover, the fact that a conventional linguistic form has a certain function does not prevent a speaker from using it to serve quite different ends. In general, there is no need that a device having a certain function be used to serve that function. A hammer can be used as a weapon, a human hair can be used as a crosshair on an instrument, and the eye-blink reflex can be used by the psychologist to demonstrate classical conditioning. Similarly, language forms are often used to serve functions that are not their own. Sometimes these extrinsic ends are ends that also interest the hearer and sometimes they are not. Uncooperative uses of conventional forms include lying to the hearer, embarrassing the hearer, insulting the hearer, purposely putting the hearer in an awkward position, and so forth. Cooperative uses include pretend uses (acting, joking) and the whole hodgepodge of Gricean implicatures (Grice 1968).

Gricean implicatures are of particular interest because they involve uses of conventional forms to produce nonconventional coordinations. Unlike cases in which, say, the speaker lies or the hearer refuses to comply, so that one frustrates the purpose of the other, a successful Gricean implicature achieves exactly the same sort of coordination that proliferates conventional language patterns. For this reason Gricean

implicatures easily become conventionalized. Speakers soon come to reproduce the relevant language forms directly on the model of previous cases of successful implicature, and hearers also reproduce the intended response that way. Hackneyed examples are 'Nice going!', 'Where's the fire?', and 'Can you . . .?', 'Could you . . .?' (as contrasted with 'Are you able to?') as used to make requests. The effect is exactly the same as when metaphors enter 'the language' (*sic!*) and become literal. The conventionalized use may continue to be associated with the original use more or less strongly, the strength of association differing also from person to person. It was only recently, for example, that I came to associate the idiom 'going haywire' with tangled hay wire, and perhaps most people, excepting farmers, never do.

For both dying metaphors and conventionalized Gricean implicatures, entrance into 'the language' is thus a matter of degree. How widely spread a conventional usage is is always a matter of degree. Language conventions can develop between just two people, as frequently occurs with twins, or develop and proliferate widely within just one generation, as in the case of certain Creoles. It does not follow that the distinction between the conventional functions of language and individual uses of language is the least bit arbitrary or trivial. Whether or not a person is dead can also, occasionally, be a matter of degree. All of the conventional functions of public language are as much its meanings as are its satisfaction-conditions. The distinction between the meanings of linguistic forms and the meanings of speakers who use them is entirely real and important.

Davidson (1986) claims that there is no boundary around the information on which a hearer may need to draw in order to interpret a speaker, hence that there is no use for the notion that a language serves as 'a portable interpreting machine set to grind out the meaning of an arbitrary utterance'. He concludes from this that not only must we 'abandon . . . the ordinary notion of a language, but we have erased the boundary between knowing a language and knowing our way around the world generally'. Chomsky is right (1992*a*, 1992*b*) that Davidson has produced here no reason to abandon the distinction between the internal systems that know language and the systems that use this knowledge. On the other hand, the systems that know language must grasp or 'represent'

(as Chomsky rather puzzlingly puts it) more than just phonetic form (PF) and logical form (LF) if the latter is understood narrowly. They must also grasp conventional function. If one draws the semantics/pragmatics distinction with the purpose of cutting between meaning and use, between 'grammatical competence' and 'pragmatic competence . . . relating intentions and purposes to the linguistic means at hand' (Chomsky 1980: 224), then the functions of all linguistic conventions fall on the semantics side of the dichotomy.

On another point Chomsky agrees with Davidson: we must 'give up the idea of a clearly defined shared structure which language users acquire and then apply to cases'. If we substitute 'structures' for 'structure', however, Davidson would surely be wrong. Language users acquire various shared structures, knowledge of public conventions, knowledge of the functions of various real-world lineages of idiom and form, without which they could not use language to communicate, nor for any of its myriad alternative uses that ride piggyback on public truth-conditional semantic conventions and conventions of 'function' in a more classical sense. There is a very simple way to grasp this point. In order to communicate with language, I must be able to predict how other people will react to my language. But no such predictions could possibly be made were it not for the possibility of conventions of use and response. These are merely repetitions of prior usage and response, including both truth-functional semantic conventions and conventions of 'function' in a more classical sense. The human language faculty plays a critical role in making it possible readily to proliferate such conventions. Whatever its history, its currently central functions are profoundly involved both with external language and with communication.

Returning to Chomsky's own views, I am not sure whether I am really one of 'his critics'. I have turned things a somewhat different way and introduced a somewhat different vocabulary to make my points. What I mean by 'public language' is not the same as Chomsky's target when he decries 'public language', and I have put a different gloss on language 'purpose' or 'function'. I have used tools he himself has fashioned at the center of my argument, claims concerning the unique capacities

of the human language faculty. There are very few points here on which I feel confident of either his agreement or his disagreement.<sup>7</sup>

## REFERENCES

- Bach, K., and Harnish, R. M. (1979), *Linguistic Communication and Speech Acts* (Cambridge, Mass.: MIT Press).
- Chomsky, N. (1980), *Rules and Representations* (New York: Columbia University Press).
- (1986), *Knowledge of Language: Its Nature, Origin, and Use* (New York: Praeger).
- (1992a), 'Language and Interpretation', in J. Earman (ed.), *Inference, Explanation and Other Frustrations* (Berkeley: University of California Press), 99–128.
- (1992b), 'Explaining Language Use', *Philosophical Topics*, 20: 205–31.
- (1993), 'Mental Constructions and Social Reality', in E. Reuland and W. Abraham (eds.), *Knowledge and Language, i. From Orwell's Problem to Plato's Problem* (Dordrecht: Kluwer), 29–58.
- (1995), 'Language and Nature', *Mind*, 104: 1–61.
- Davidson, D. (1986), 'A Nice Derangement Of Epitaphs', in E. Lepore (ed.), *Truth and Interpretation* (Oxford: Blackwell).
- Dawkins, R. (1976, rev. 1989), *The Selfish Gene* (Oxford: Oxford University Press).
- Gilbert, M. (1983), 'Agreements, Conventions, and Language', *Synthese*, 54: 375–407.
- (1989/92), *On Social Facts* (London/NY: Routledge; Princeton, NJ: Princeton University Press).
- Grice, H. P. (1968), 'Logic and Conversation', in Grice, *Studies in the Way of Words* (Cambridge, Mass.: Harvard University Press).
- Kaplan, D. (1990), 'Words', *Aristotelian Society Proceedings, Supplementary Volume*, pp. 103–19.

<sup>7</sup> I am enormously grateful to Louise Antony for so patiently, thoroughly, and insistently misunderstanding an earlier version of this essay, forcing me, kicking and screaming, into actually saying what I meant in many places. Indeed, in many places she even forced me into knowing what I meant. Without question, she is the most brilliantly dense reader I have ever had and I hope and think that the essay was immeasurably improved by her interference.

- Lewis, D. (1969), *Convention: A Philosophical Study* (Cambridge, Mass.: Harvard University Press).
- (1975), 'Languages and Language', in K. Gunderson (ed.), *Language, Mind and Knowledge* (Minneapolis, Minn.: University of Minnesota Press), repr. in H. Geirsson and M. Losonsky (eds.), *Readings in Language and Mind* (Oxford: Blackwell, 1996), 134–55.
- Lieberman, A. M. (1999), 'The Reading Researcher and the Reading Teacher Need the Right Theory of Speech', *Scientific Studies of Reading*, 3/2: 95–111
- Millikan, R. G. (1984), *Language, Thought, and Other Biological Categories* (Cambridge, Mass.: MIT Press).
- (1993), *White Queen Psychology and Other Essays for Alice* (Cambridge, Mass.: MIT Press).
- (1997), 'Images of Identity: In Search of Modes of Presentation', *Mind* 106/423: 499–519.
- (1998), 'A Common Structure for Concepts of Individuals, Stuffs, and Basic Kinds: More Mama, More Milk and More Mouse', *Behavioral and Brain Sciences*, 22/1 (Feb.), 55–65; repr. in E. Margolis and S. Laurence (eds.), *Concepts: Core Readings* (Cambridge, Mass.: MIT Press).
- Pinker, S. (1994), *The Language Instinct* (New York: William Morrow).
- Recanati, F. (1987), *Meaning and Force: The Pragmatics of Performative Utterances* (Cambridge: Cambridge University Press).
- Schiffer, S. (1972), *Meaning* (Oxford: Clarendon).
- Searle, J. (1969), *Speech Acts: An Essay in the Philosophy of Language* (Cambridge: Cambridge University Press).
- Sellars, W. (1963), *Science, Perception and Reality* (New York: Humanities).
- Strawson, P. F. (1974), *Subject and Predicate in Logic and Grammar* (London: Methuen).

### 3 On Meaning, Meaning, and Meaning\*

---

#### *Introduction*

Various crisscrossing distinctions have been drawn in the philosophical tradition between kinds or dimensions of linguistic meaning or between meaning and other dimensions of linguistic function. In this chapter I'll try to collect together from various books and papers the results of my own investigations on different aspects of meaning. The underlying idea is that to understand how language works one must look, first, to the cooperative functions that various language forms perform, understanding these on a biological model as what these forms accomplish that keeps them in circulation. To explain the cooperative function of a language form is to explain its survival value, the source of its proliferation, what it does that accounts at the same time for the fact that speakers continue to use it and that hearers continue to react to it often enough in standard ways. Next we should look at language mechanics, at *how* language forms perform their functions. For some language forms there are conditions in the world that are necessary to support their functions and that vary systematically with certain variations in the forms themselves. These are truth-conditions, and they are determined by a kind of 'meaning' that I will call 'semantic-mapping functions'—'functions', this time, in the mathematical sense. (Semantic-mapping functions determine truth-conditions; truth-conditions only delimit and do not determine semantic-mapping functions. I will get to this.) Last we need to describe the psychological mechanisms that are involved in implementing the functions of various language forms: the ways that

\* © Ruth Garrett Millikan 2005.



speakers and hearers manage to produce and understand these forms so as to promote performance of their cooperative linguistic functions.

There is also an important distinction, of course, between speaker meaning and linguistic or conventional meaning. This is the difference between the cooperative function, with its associated truth-condition and so forth, of a public-language form and functions that individual speakers may use or try to use the form to serve. I will discuss this distinction, but only to set speaker meaning aside. My proposal is that there are these three basic kinds of linguistic meaning:

- (1) conventional linguistic cooperative functions, to be called ‘stabilizing functions’;
- (2) conventional semantic-mapping functions (‘functions’ in the mathematical sense) which determine truth- and other kinds of satisfaction-conditions;
- (3) methods of identifying—to be called ‘conceptions’ and ‘conceptual components’—that govern individual speakers’ grasps of referents and of truth- or satisfaction-conditions, hence help to determine their dispositions to use and understand various conventional language forms.

I will argue that the third of these, conceptions governing individual speakers’ grasps of referents and satisfaction-conditions, may exhibit little or no overlap among competent members of the same language community. Thus, none of the aspects of meaning that I will define corresponds at all well either to any traditional notion of intension or to any Frege-related notion of sense.<sup>1</sup> The meanings that characterize the public part of a language are fully extensional. I will have to say quite a lot about that before I am finished. Let us begin with simpler matters.

## 1. *Stabilizing Functions*

When speakers are conforming to the conventions of a language, speaker meanings coincide with conventional meanings. But speakers

<sup>1</sup> This despite the embarrassing fact that in my (1984) ‘(Fregean) sense’ was the name I gave to what I now call ‘semantic mapping’ and ‘intension’ my name for what I now call ‘conception.’ I had my reasons, but they were not good.

often use language forms for purposes that diverge from conventional meanings. To distinguish speaker meanings from conventional linguistic meanings we need to know what a public-language convention is. Probably the best known theory of convention and its application to language is that of David Lewis (1969, 1975). I will clarify my position by comparing it with his.

Lewis describes a convention as a regularity in the behavior of a population such that within the population there is mutual knowledge (1) that everyone conforms to the regularity, (2) that everyone prefers to conform given that the others do, and (3) that everyone expects everyone else to conform for the same reason he does. The reason each prefers to conform is that conforming solves a coordination problem. A coordination problem arises when people have a purpose in common which must be achieved by joint action, where the contribution that each must make will vary depending on what each of the others contributes, and where there is more than one acceptable way of combining contributions to produce a successful outcome. Then, coordination is necessary. It is best for everyone if everyone makes his contribution according to the same solution plan. To each it doesn't matter as much which plan is chosen as it matters that the same plan is chosen by all. In many cases, Lewis says, the plan that is chosen will be the one for which there is a precedent. It has been used before, which makes it a salient plan—one that comes to mind and that each participant assumes will come to the mind of the other participants. Each participant thus steps into his role in this plan on the assumption that the others will adopt their roles according to the same plan. When a precedent for solving a coordination problem spreads in this way a convention is born. Thus, Lewis claimed that social conventions of all kinds, including linguistic conventions, are supported by rational beliefs and intentions concerning one another's thoughts. It is true that children and idiots may conform to the conventions of language without having reasons of this sort but, Lewis claims, 'they are not parties to the convention and their linguistic competence is incomplete' (1969: 51).

Now, I agree that the conventions of language arise and spread because they solve certain kinds of coordination problems. Not all conventions solve coordination problems, however. And for those that spread because they do, the 'because' is almost never a reasoned because but some more

mundane kind of *causal* because. The rest of us conform to linguistic conventions in exactly the same unreasoned way that the idiot and the child do. Further, despite apparent consensus among philosophers that conventions always involve regularities of behavior within a group, my claim is that conventions do not generally require regularities of behavior, either *de facto* or *de jure*. In particular, conventional *coordinations*, including linguistic coordinations, do not, in general, require regularities of behavior. These claims were defended in Chapter 1. Here I will just review the high points.

A convention, in the sense that a natural language contains conventions, is merely a pattern of behavior that is (1) handed down from one person, pair, or group of persons to others—the pattern is reproduced—and (2) is such that *if* the pattern has a function, then it is not the only pattern that might have served that function about as well. Thus, if a different precedent had been set instead, a different pattern of behavior would probably have been handed down instead. Putting a wreath on the door at Christmas time, dyeing eggs for Easter, and drinking green beer on St Patrick's Day are conventions in this sense. In Japan the convention is to eat with chopsticks, in America, with a knife and fork. Against Lewis, that these are conventions (1) does not necessarily mean that they solve coordination problems. Also, (2) it does not necessarily mean that they are universally followed. Indeed, there are many conventions for which conformity is neither prescribed nor mandatory in any sense. Of course, some conventions, such as driving on the right in the United States, do solve coordination problems, are universally followed, and are mandatory. But that is not what makes them conventions. Also, linguistic conventions do solve coordination problems, but they are neither universally followed nor mandatory.

When a conventional pattern of behavior is handed down because it is solving a coordination problem, the mechanism for this is usually quite simple. No matter how the precedent for the convention was originally set, if the coordination it effects is an obvious and important one it will tend to proliferate without anyone's thinking about anyone else's thoughts. Like other higher animals, people repeat behaviors

that have been successful in achieving wanted results in the past. Unlike most other animals, they tend also to copy behaviors of others that have been successful in producing wanted results. Behaviors that constitute solutions to coordination problems achieve results desired by all parties to the coordination, hence these behaviors will tend to be reproduced when similar results are desired. There is no need for the various parties in the coordination even to recognize the problem as a coordination problem, let alone to think about one another's thoughts in order for the convention to proliferate. If other people are driving on the right, then I will drive safely only if I drive on the right. Thus, I might learn to drive on the right without ever quite realizing that it is only a convention to do so. Exactly thus not only children but very smart primitives typically are unaware that the languages they speak are merely conventional.

Specific language forms continue to be reproduced by speakers within a language community merely because, often enough, they prompt hearer responses that contribute to the fulfillment of speaker purposes in speaking. Similarly, hearers continue to respond in conventional ways, for example by believing or by doing what they are told, because, often enough, the result is rewarding for them. Often enough, believing or doing what one is told leads to believing or doing what is useful or what will keep one out of trouble. Speakers within a language community are, simply, *adapted* to an environment in which hearers are responding, sufficiently often, to the forms speakers produce in ways that reinforce these speaker productions. Correlatively, hearers in the community are, simply, *adapted* to conditions under which speakers, sufficiently often, produce these language forms in circumstances such that making conventional responses to them aids those hearers.

Consider, for example, a speaker whose purposes in using the word 'dog' will be achieved only through calling attention to dogs or to facts that concern dogs or through changing hearers' behaviors toward dogs. Such a speaker will eventually stop trying to use the word 'dog' for these purposes if they are never achieved. Also, a hearer whose language-understanding faculties turn his mind to dogs whenever

speakers use the word 'dog' will soon unlearn this response if speakers never use the word 'dog' such that it carries information or expresses intentions that concern dogs. Similarly, consider those syntactic forms that get labeled 'indicative' in various languages. These forms usually have a number of alternative functions, but no form will be so labeled unless one of its functions is to effect production of true beliefs having propositional contents carried by other aspects of these sentences. These conventional forms are surviving in part because, often enough, this particular effect is of interest both to speakers and to hearers. Production of false hearer beliefs may occasionally interest speakers, but rarely serves the interest of hearers. A hearer unable to interpret the indicative sentences he hears so as sometimes to extract genuine information from them would soon cease to form beliefs on their basis. And if hearers ceased ever using indicative sentences as guides in forming beliefs, speakers would stop trying to use them for purposes that required imparting beliefs. Similarly, if it were not sometimes in the interest of hearers to comply with imperatives—advice, instructions, directions, friendly requests, sanctioned directives, and so forth—they would soon cease ever to comply. And if hearers never complied with imperatives, speakers would soon cease to issue them. Imperative syntactic forms would become obsolete.

A corollary is that the functions of public-language forms are not on the same level as either speaker purposes or hearer purposes taken alone. The conventional functions of language forms are not, for example, merely standard speaker purposes. Conventional language forms are selected for performing services satisfactory at once to both partners in communication. Their functions must balance speaker with hearer interests. Because the conventional function of a linguistic form will remain stable only if it continues to serve the interests of both speakers and hearers often enough, I call it a 'stabilizing function'. Linguistic 'meaning' in the sense of stabilizing function is on an entirely different level from, for example, average speaker meaning.

Similarly, on this analysis a linguistic convention consists in a pattern that includes *both* a conventional contribution by the speaker *and* a conventional contribution by the hearer. The hearer's contribution is as much a part of the convention as is the speaker's. Thus, the linguistic

convention includes important aspects of what Austin called the *perlocutionary* act. As such it effects a genuine coordination between speaker and hearer, each of whom must play his part if the coordination is to be successful. Contrary to this, Lewis claimed that '[a] member of the audience, as such, is not constrained by convention . . . Only when he takes his turn as communicator does he himself act in conformity to the convention of truthfulness in *L'* (1969: 179–80).

We can see why Lewis took this position. According to his analysis, a convention was a regularity of behavior in a population such that everyone conforms to the regularity (1969: 42, 58). Later, he modified this to allow that *almost* everyone conform *almost* all of the time (1969: 78), but still, it is plainly false that almost every hearer of a directive complies with it and plainly false that almost every hearer of a description believes it. So on Lewis's account the hearer's response could not be part of a linguistic convention.<sup>2</sup> Lewis also says 'forming a belief . . . is normally not a voluntary action and hence not an action in conformity with convention' and, of directives, '[e]ven if the audience should act, the action may not answer to an interest common to the communicator and the audience' (p. 180). But if, as I have claimed, a convention is merely a reproduced pattern whose form is arbitrary with respect to its function, then there is no requirement on how voluntarily or how regularly the pattern is reproduced or on how often the pattern is broken, with either speaker or hearer failing to contribute his or her proper part (Ch. 1; Millikan 1984: ch. 4). Sometimes the speaker is not interested in genuine cooperation. Sometimes the hearer is not. Sometimes mistakes are made. Conventional coordination patterns need to succeed only often enough to avoid extinction.

Notice as well that alternative coordination conventions serving the same purpose often happily exist side by side in a community. Besides linguistic conventions, Lewis talked of 'signaling conventions', which he did describe as involving the receiver's responses as well as the signaler's gestures. He illustrated with signals used by a man standing outside to help a truck driver back into a tight space. The helper and the driver both

<sup>2</sup> Lewis also claimed however that the conventions of a language involve speaking the truth, which may be just as implausible on his assumption that conventions require nearly universal compliance.

want to maneuver the truck safely into the space. How the signals are composed by the helper and how they are read by the driver have to be coordinated if this common end is to be achieved. Lewis apparently overlooked that both helper and driver might easily be acquainted with more than one signaling system commonly used for this purpose. So long as these different systems didn't happen to contain identical signals that meant different things, it wouldn't matter which system the helper chose to use, the driver would recognize the signal and follow it. What is necessary for success is only that the same precedent should be followed by both helper and driver on each individual occasion. What they do on other occasions doesn't matter. An initiating move by one party will immediately be recognized by the other as coming from a particular lineage of precedent with which both are familiar. Thus it is that linguistic conventions are neither universally followed nor mandatory. Many alternative conventions can possess the same stabilizing function.

Language conventions can be considered as lineages of precedent. A public language is a huge web of crisscrossing lineages of reproduced patterns consisting of tokens of linguistic forms and responses to them. People listen to one another, then repeat words and idioms they have heard, syntactic arrangements they have heard, and tonal inflections they have heard, arranging these into new combinations. Words, idioms, syntactic forms, tonal inflections, and so forth are handed down from one person to others because these elements are helping to serve coordinating functions. These stabilizing functions are, in one of that term's various senses, their 'meanings'—the first of the three 'meanings' listed in this essay's title. One thing to investigate then is exactly what kinds of stabilizing functions compose these meanings. What various jobs do linguistic forms do to keep themselves in circulation?

The thesis that linguistic conventions correspond to reproducing lineages of cooperatively used tokens-with-responses has direct implications for the individuation of linguistic forms. For the purposes of semantics, what makes two tokens be tokens of the same linguistic type is not their sound or shape, or the phonemes or letters of which they are composed, or their surface syntactic arrangements. They are tokens of

the same type only if they have been copied from the same pool of tokens reproducing in the same language community. They must be segments from the same historical lineage. Genuine words cannot be accidentally formed by the wind. Further, any genuine linguistic token is automatically a token from one particular language or another. When discussing linguistic forms, reference to the form as being 'in  $L_1$ ' or 'in  $L_2$ ', etc. may help the hearer to identify the form intended, but whether identified or not, if the form is a genuine natural language form it already is *essentially* either in  $L_1$  or  $L_2$  or some other language without that. Otherwise it is not a linguistic form, but merely a describable shape or sound.

Unlike the lineages that make up animal species, linguistic lineages frequently acquire new functions without changing their physical forms. Similar to mutations in biological evolution are novel uses of conventional linguistic forms introduced by speakers through figures of speech or through Gricean implicature. If the hearer understands the figure or the implicature, the novel use will serve a new coordinating function. It may then be copied by other speakers and may in time be understood directly by hearers without having to go through the process of unpacking a figure or an implicature. Then a new lineage of tokens with a different stabilizing function has branched off from the original lineage but without any change in physical form. Suppose, for example, that a new metaphor is copied again and again. For a very long time those who use it and those who understand it may continue to read it as originally derived from its original source. For most speakers and hearers, acquaintance with the old lineage and with the new lineage may both together and equally be responsible for its use and for its easy comprehension. Later, however, the new use may become as familiar as the old and may start to be proliferated quite independently. Then the metaphor becomes 'dead'. An entirely independent branch of the family has been formed from tokens with exactly the same physical form. The result is called 'polysemy'—one sound, many meanings.

Families of linguistic forms quite typically form wide-spreading bushes, many different branches having slowly formed over time, and more branches from those branches. Since branches often take a long time to separate off completely, the places where true branchings begin



are not at all sharp. At a given time, exactly how many branches there are is not definite. Still, it is helpful to give a name to the branches of a given form that are currently fairly independent, each being well enough established that it would survive even if all the others should die out. These independent branches I call 'least types'. Least types correspond to the various relatively independent stabilizing functions (different 'senses') of a polysemous language form (Millikan 1984: ch. 4). However, typically the number of such 'meanings' and the divisions between them are very far from precise.

Surface syntactic forms may also branch into independent least types. They too may be polysemous, having a number of branches that propagate more or less independently. For example, definite descriptions are sometimes used by speakers merely to identify a particular referent for a hearer, the description itself being of no interest at all. At other times definite descriptions are of interest in their own right; indeed, the speaker may not care whether the hearer identifies the referent or not so long as the description is remembered. Donnellan's distinction (1966), when understood this way, is a distinction between two stabilizing functions of definite descriptions that tend to divide these descriptions into two least types. Definite descriptions tend to be polysemous in stabilizing function.<sup>3</sup> Similarly, consider the various grammatical moods in a language. Sometimes you impart beliefs with the indicative mood but at other times you use it to give orders. Sometimes you ask questions with the interrogative mood but at other times you use it to make requests. A number of relatively separate conventions are helping to propagate the same surface syntactic forms. Differences in what Austin called 'illocutionary' and also 'perlocutionary' force may often be carried conventionally by syntactic forms that are polysemous in this way (Millikan 1984: ch. 4; Ch. 10). Then, understanding the force of a particular utterance using the form will require disambiguation

<sup>3</sup> It doesn't follow that they are polysemous in semantic-mapping function or that they correspond to more than one kind of truth-condition. I will return to this soon.

Donnellan's claim that for one use of definite description it doesn't matter whether the description is empty so long as the speaker's intended referent is understood by the hearer is best interpreted as a claim about speaker meaning rather than linguistic meaning. Certainly, it is not a stabilizing function of definite descriptions to bring things not correctly described by them to hearer's minds.

from context, just as understanding the meanings of individual words often does.

## 2. *Semantic-mapping Functions (and Satisfaction-conditions)*

What sets linguistic acts and, more generally, communicative acts apart from other acts with cooperative functions is that communicative forms work in part by mapping or, as Wittgenstein put it, 'picturing'. They correspond to states of affairs in accordance with semantic-mapping functions that have been determined by convention.<sup>4</sup> Directive communicative forms have as their stabilizing functions to yield states of affairs—completed actions—that vary with variations in the sentences exhibiting these forms. For example, directive least types used in giving orders have as stabilizing functions to produce compliance, what constitutes compliance being determined along the lines Tarski proposed. The state of affairs that would result from compliance is the 'satisfaction-condition' of the directive sentence. Descriptive communicative forms have stabilizing functions that can be performed through normal mechanisms only if they correspond to states of affairs existing independently. For example, conventional fact-stating least types are designed to produce true beliefs in hearers, but a true belief will be formed by normal mechanisms only if the sentence corresponds to a world affair in accordance with its conventional semantic-mapping function. False sentences do not cause true beliefs in hearers through normal mechanisms. The truth-conditions of a descriptive sentence are also called 'satisfaction-conditions'.

The semantic-mapping function for a sentence determines the sentence's satisfaction-condition, but the satisfaction-condition does not determine the semantic-mapping function. The semantic-mapping function is given by rules according to which significant transformations of the sentence that conserve its syntactic form yield different truth- or satisfaction-conditions. Compare the sentence 'It's raining'

<sup>4</sup> In the case of most animal communication determined genetically.

with the sentence 'Rain is falling here now'. 'It's raining' contrasts with 'It's snowing', 'It's hailing', 'It's sleeting', and so forth. All display the same syntactic form, the transformations that substitute 'snow', 'hail', and 'sleet' for 'rain' determining different satisfaction-conditions in a systematic way. Similarly, 'Rain is falling here now' contrasts with 'Snow is falling here now', 'Hail is falling here now', 'Sleet is falling here now', and so forth, but it contrasts, further, with 'Mist is rising here now', and with 'Rain was falling in Rome yesterday'. The truth-conditions of 'It's raining' and of 'Rain is falling here now' are the same, but the semantic mapping is different. 'Many drops of water are presently precipitating from the atmosphere and landing close to this place' also has the same truth-condition but is articulated by yet another semantic-mapping function. For vividness, compare the semantic-mapping function for a bee dance with that of an English sentence having the same truth-condition. Bee dances show by the angle of their axis where there is nectar relative to a line between the hive and the sun, but there are no transformations of the bee dance that would tell about nectar location relative to objects other than the hive and the sun, or about the location of anything other than nectar. Only reference to the angle between the nectar and the line from the hive to the sun can be varied in the bee dance. Further, the bee dance is not subject to a negation transformation. No English sentence with the same truth-conditions approaches this degree of inarticulateness. The semantic mapping of a sentence articulates it, placing it in a logical space of contrasting possibilities. Its truth-condition is not, as such, articulated.

Stabilizing functions can vary while semantic mapping remains the same. Compare 'Jane will close the door' to 'Will Jane close the door?' And there are more interesting examples. Wilfrid Sellars claimed that the function of the sentence form ' "X" means Y', as in ' "Rot" means red' and ' "und" means and', is to produce in the hearer a disposition to use 'X' in the same way he already knows to use 'Y'. The 'Y' in this rubric, Sellars said, is neither mentioned nor used in the usual way (Sellars 1956). It is used in a special way—held up, as it were, as a model (Millikan 2004: ch. 7). Compare the function of the form ' "X" and "Y" are used the same way'. Here 'Y' is mentioned rather than used. This

sentence has the same truth condition as ' "X" means Y' but its function is different. Its function is to produce a belief about words, whereas the function of ' "X" means Y' can be performed even if the hearer lacks a concept of words (as very young children apparently do—Susan Carey, private correspondence).

Peter Strawson claimed that the function of the identity form 'A is B', as in 'Cicero is Tully', is to induce the hearer to merge all of the information he has accumulated under the concept he associates with the word 'A' with the information he has accumulated under the concept he associates with 'B', so that he no longer harbors this information under two separate concepts (Strawson 1974). More accurately, the stabilizing function must be to induce the hearer to do this appropriately, such that the resulting concept is not confused or equivocal (Millikan 2000). If this is the function of the form 'A is B', then its truth-condition is the same as for the form ' "A" has the same referent as "B" ', in which the 'A' and the 'B' are mentioned rather than used. But these two sentence forms do not have the same function. The hearer of ' "A" has the same referent as "B" ' is to form a belief about words, hence needs concepts of words and also the concept of reference, whereas the hearer of 'A is B' needs neither.

I have claimed (Millikan 1984: ch. 12) that the stabilizing function of the form 'A does not exist' is correctly to induce the hearer to disengage his concept associated with 'A' from ordinary referential uses, relegating it, for example, to pretend uses, or eliminating it entirely from his conceptual repertoire. Correlatively, the function of 'A exists' is correctly to engage a previously disengaged concept associated with 'A'. But if these are their functions, the sentence forms 'A does not exist' and 'A exists' have the same truth-conditions as do ' "A" has no referent' and ' "A" has a referent'—though, again, the functions of these sentences are not the same. The latter have as stabilizing functions to cause beliefs about words.

Adding a different kind of example, the two uses of definite descriptions mentioned above in connection with Donnellan's distinction may correspond to two independent stabilizing functions of these, but these two uses require exactly the same conditions for truth. The world affair

needed to make one least type of definite description serve its stabilizing function through normal mechanisms is exactly the same as that needed for the other. In both cases the truth-condition is Russellian.<sup>5</sup>

The study of semantic-mapping functions should include a study of the peculiarities of indexicals and demonstratives. As I understand it, there is more than one way of describing their semantic-mapping functions and I have suggested somewhat different though, I believe, compatible ways of thinking about the matter in my (1984: ch. 10) and (2004: ch. 12). I omit discussion of these forms here, my purpose being only to make clear why we need to distinguish among the three broad aspects of meaning mentioned at the start.

### 3. *Conceptions*

I think that Frege made a mistake in positing something common beyond *Bedeutung* that is grasped by the mind of every competent speaker using the same unambiguous linguistic form. A related mistake suffuses the tradition of conceptual analysis in seeking shared 'criteria' for the correct application of various terms, criteria taken to be learned, in some mysterious way, when one learns one's language. On the contrary, the public meaning of a simple referential term typically includes only its stabilizing function and its reference, and since the stabilizing function depends almost entirely on sentential context, the public meaning is essentially *just* reference.<sup>6</sup> I intend this sweeping assertion to include terms for properties, kinds, stuffs, and so forth, which I will treat here as also being, in a broad sense, referential. The claim will need qualifications, but first I'll just try to explain it.

The idea to be opposed is that for different users to understand the same referential term as having the same meaning requires that their psychological processing be similar in certain ways. The idea to be defended is that 'Meaning is not in the head'. But I want to launch an

<sup>5</sup> See n. 4 above.

<sup>6</sup> See (Millikan 1984: ch. 4) on the most general stabilizing function of all referential terms.

attack on Frege's idea and on conceptual analysis that is more radical and exhaustive than the familiar offenses once launched by Putnam, Burge, and Evans. My argument grows out of a view about thought structure, a view about what it is to have a concept of a property, an individual, a kind or a stuff, and so forth. It is fully articulated in Millikan (1984, 2000). Here I can make only a small sketch.

Consider what is involved in being able to recognize, for example, shapes. Think of the variety of proximal visual stimulations to which a given shape may give rise when viewed from various angles, from different distances, under different lighting conditions, through various media such as water or fog, when colored different ways, when partially occluded, and so forth. How shape constancy is achieved by the visual system, the capacity to recognize the same shape as the same under a range of conditions, is a problem of nearly unimaginable complexity that is still largely unsolved.

Similarly, how color constancy, texture constancy, size constancy, and distance constancy are achieved are enormously complicated problems. We are also adept at perceiving sounds, especially speech sounds, as the same sound at origin whether near or far, through air or through water, muffled or distorted, and so forth. What does seem clear in each of these cases is that no single rule is applied. Different clues are used by the perceptual systems in different circumstances, separately or together. For example, distance is perceived with the help at least of ocular disparity, tension in the focusing muscles, occlusion of one object by another, knowledge of the size of objects viewed, and atmospheric haze. We also recognize distances by touch and stretch using many different parts of the body, and we recognize distances of things that make noises fairly well by ear. And of course there are more readily noticed ways of recognizing distances as well; for example, by measuring with a ruler or a tape measure or just a string, or measuring as a surveyor does by triangulation, or measuring with an odometer or a micrometer, or by the time of the return of light. None of these ways of telling distances is infallible, nor is any definitional of our concepts of distances. On the other hand, each adds something to our concepts of distances, nor could we have distance concepts at all were we not in command at least of some of these methods of recognition.

The situation is similar, if not always so extreme, with our grasp of other perceptual constancies. The perceptual systems do their work in flagrant violation of the ideal once set by champions of operational definitions. The more ways the better when it comes to methods of perceiving a property. After all, the ways in which empirical properties affect the various senses through intervening media is a thoroughly empirical matter, a question of natural law, not a matter of logic or definition. That is why neither phenomenalism nor verificationism could ultimately survive.

Now, it is conceivable that all normal persons perceive some constancies, for example depth, in the same way, conceivable even that they are genetically programmed rather than perceptually tuned to perceive some of these constancies in standard ways. The issue is under debate. But surely whether one's perceptual capacities were entirely normal in this regard would not affect what one meant by the English words one used in designating depths or shapes or textures. Being blind in one eye so that one could not perceive depth using ocular disparity would not change what one meant by 'near' and 'far', nor, indeed, is it sensible to suppose that Helen Keller meant something different by 'near' and 'far' than you do. I don't want to debate about whether there are secondary qualities,<sup>7</sup> but surely we recognize in perception enough properties and relations that obviously are primary to make the point. There are many different ways of recognizing each of these properties, but none *defines* either the property or the words that stand for it.

Turning now to the opposite extreme, consider proper names. Besides having a referent does your name have a definition? What is involved in someone's understanding who's meant by your name—say, a child in your family, your child's teacher, a student of yours, the student's wife, a reader of your essays, the pharmacist who fills your prescriptions. Do these people all understand who's meant by your name in the same way? The reasonable answer is that there is no special thing common in the minds of all people who understand your name except, I have argued (Millikan 1984: chs. 4, 9; 2000: ch. 6), some practical capacity to

<sup>7</sup> The very fact that psychologists can make a study of how color constancy is perceived seems to cast doubt on the idea that colors are secondary qualities, at least in the Lockean sense of that term.

reidentify as such, in actual context, the *least type* that is *your* name (rather than the name of someone else with 'the same name'), so as to recognize when information is being offered about the same person, you, again. Speaking more generally, what it is to have a concept of an individual is, in part, to have an ability to recognize, in one way or another, under at least some circumstances, when one is encountering information concerning that individual, and one recognizable way that one encounters information about a thing, besides through direct perception, is by encountering sentences that contain its name. To defend this position properly the right characterization of information is needed (Millikan 2004: chs. 3–4) and the right story about perception through language (Millikan 1984: ch. 9; 2000: ch. 6; 2004: ch. 9), as well as a story about abilities that allows for their fallibility (Millikan 2000: ch. 4). But that names of individuals need not be associated either with independent publicly agreed on ways of recognizing these individuals in order to do their work or with agreed on descriptions associated with these individuals is generally accepted, I think.

Call the sum of the various ways that you have of recognizing a thing or, what amounts to the same, of recognizing when you are receiving information about a thing your 'conception' of that thing. Your conceptions of most common things have many components, for you have many ways of recognizing these things—no infallible ways, of course, but many fairly reliable ways. Whatever you know about a thing is part of your conception of it too, for whatever you know might help you to identify it, or help prevent you from misidentifying it, under some circumstances. Some components of conceptions are explicit, involving the use of descriptions, hence of prior concepts, in their application. Other conceptual components are implicit, moving one directly from perceptual experience to an identification of what is perceived. My claim so far is that neither the names of perceivable properties nor the names of individuals are associated with conceptions or conceptual components, either explicit or implicit, that all users of their names must possess in order correctly to understand these names. No specific way of identifying their referents is required. True, in some cases there does exist considerable overlap in the conceptions that most people use in understanding the referent of such a name; for example, since 'Mark



Twain' was Samuel Clemens's pen name, a large proportion of people know that Mark Twain was a writer, perhaps even that he wrote *Huckleberry Finn*. You can usually count on someone having that knowledge if the name 'Mark Twain' is in their vocabulary. And the implicit conceptual components by which we recognize many common properties may be shared among a large majority of adults. But if someone were born with bat's ears and could only hear shapes, this would not prevent him from learning the English words 'round' and 'square'. I take it that this much is not highly controversial. The principle can be extended, however, to less obvious cases.

Many terms for kinds name kinds that are objective natural units, discovered rather than created by thought and language (Ch. 6; 1984: ch. 16–17; 2000: ch. 2). These 'real' kinds are important subjects for knowledge because there is a reason why the various members of the kind mostly resemble one another in a good number of ways, hence there is a reason why one can learn from observation of one or a few examples of the kind much that is likely to be true of other members. Most single terms designating kinds designate real kinds of this sort (Ch. 6; 2000: ch. 3). Typically, these kinds not only have many properties, there are also many ways to recognize them. Think how many ways there are of telling that something is copper, or that a dog is present. Do you have to look to tell it's a lemon? Or that it's raining? How much of what portion of 'The First Noel' or the Lord's Prayer do you have to hear to recognize it? To have a valid concept of a real kind one needn't know the reason for the resemblance of its members, what natural principles hold the kind together. One only needs some fairly reliable ways of reidentifying the kind—the more the better, of course, since most ways are only applicable on some occasions. Like concepts of individuals, concepts of real kinds can be supported by alternative conceptions, alternative methods of recognition, and there are no conceptual components that all users of a real kind's name must possess to understand it (Ch. 6; 1984: ch. 9; 2000: chs. 3, 5).

The third aspect of meaning, conception, is not then essentially public. It attaches in the first instance to idiolects rather than public languages. However, there usually is considerable overlap among people's conceptions corresponding to names of very common real kinds. Also,

sometimes conceptual components are passed on explicitly from generation to generation; for example, the definitions of certain geometrical figures. One *could*, after all, 'define' a circle, instead, as a closed plane figure with but one side of uniform curvature, but it is not conventional to do so. And in the case of fictional names, and empty names like 'phlogiston' and 'witch' when these are mistakenly thought to have referents, there is no public meaning beyond certain traditional explicit conceptual components, traditional descriptions, passed down from person to person. There is no more to public meaning in these cases than public conception; indeed, public conception that is highly subject to drift. Santa Claus acquired red and white attire and reindeer rather late in his career, while phlogiston and witches took on different diagnostic properties over time in the eyes of different investigators. Water and dogs, by contrast, are surely recognized by us, practically all of the time, in exactly the same ways they were by the ancients.

Traditional descriptions associated with empty terms fail to reach anything real, hence do not correspond to real abilities to identify. But having empirical concepts, having thoughts of objects, properties, and so forth, essentially involves abilities to identify. It seems to follow that empty terms do not express real concepts. This brings us to the externalist core of this chapter on meaning.

The claim is that the meaning of an empirical term is, in the first instance, its referring to something, and only in the second instance ways one has of identifying this thing through various of its manifestations. Wittgenstein was right, after all, that the primary check on whether we mean the same by our words is agreement in judgments, but agreement in judgments proves nothing about agreement in the methods of identifying used in making those judgments. It is clear that nothing inside the head or mind can determine, in and of itself, whether one's dispositions to react to sensory stimulations with would-be thoughts of individuals, properties, and kinds manifest real abilities to identify such things or not. Similarly, nothing inside the head determines whether those explicit inference-dispositions whose job is to help one identify such things are actually doing their job. But if a would-be thought has conceptual components that are explicit, and if the prior concepts in the descriptions employed in these components are not

themselves empty, then there is a legitimate, though secondary, sense in which even a term expressing an empty thought can have a meaning—because components of its conception have meanings. Indeed, if the term is public, it will have conceptual components that are both explicit and traditional, having been handed down from speaker to speaker, hence it will have a sort of *public* meaning. But suppose there were an empty concept that had only implicit conceptual components, that was not anchored by any prior nonempty terms. Such a term would have no more claim to membership in the realm of the intentional or the semantic than a sneeze. It would merely be a quirkish regular response to certain sensory stimulations, resulting, presumably, from the faulty operation of systems *designed* to *design* genuine concepts, genuine thoughts, through experience, but that had failed in that task.

A crucial task incumbent on any advocate of meaning externalism is to explain how we acquire evidence through experience that our concepts are not empty, that they are anchored externally in what is objectively real. The externalist is obliged to accompany claims about the ontology of meaning with a plausible *epistemology* of adequacy for empirical concepts. She must construct an epistemology of meaning to support her claims in the philosophy of mind. I consider this an urgent matter, though one sorely neglected in the current literature on externalism. The epistemology of concepts, or of meaning, is the subject of Millikan (1984: chs. 18–19; 2000: ch. 7; 2004: ch. 19).

#### 4. *Replacing Intensions and Fregean Senses*

As said at the start, none of my trio of meanings corresponds at all well either to any traditional notion of intension or to any Frege-like notion of sense. Both these latter notions were introduced on the assumption that a grasp of certain ways of identifying or certain properties by which a thing may be identified must be shared by users of any public term that refers to it, whereas I claim that any such grasp is, in the first instance, a private matter. But this claim needs to be defended with an alternative explanation of the phenomena that lead to the postulation of

intensions and Fregean senses. I take it that there are three central classical arguments for something like intensions or senses, one from the informativeness of sentences asserting identity, a second from the need to analyze statements asserting existence, the third from the behavior of referential terms in intensional contexts. I will briefly discuss each in turn.<sup>8</sup>

First, identity statements. I have agreed with Strawson that the stabilizing function of an identity sentence 'A is B' is to encourage the hearer to merge under a single concept all of the information she has accumulated under the concept she associates with the word 'A' with that under the concept for 'B'. More precisely, it serves its stabilizing function by joining the conception the hearer has associated with 'A' to the conception associated with 'B' so that these now (correctly) govern the same concept. Rather than inducing *beliefs*—compare beliefs to mental sentences—it alters *conceptions*, ways of identifying (Millikan 2000: ch. 12). Thus, for any hearer who associates a different conception with 'A' than with 'B', the effect of a true identity statement 'A is B' obviously is different from that of 'A is A'. This can be true and important even if no two hearers who react in the stabilizing way to 'A is B' happen to share their conceptions associated with 'A' or with 'B'.

Second, existence statements. Statements of the form 'A doesn't exist', I have claimed, induce a hearer to disengage his concept associated with 'A' from ordinary referential uses, eliminating it entirely or reserving it only for pretend uses. Statements asserting the existence of A reverse this effect. To engage or disengage a concept is the same as to engage or disengage the conception that governs that concept. The forms 'A exists' and 'A doesn't exist' can serve these functions regardless of how diverse listeners are in the conceptions they associate with the name 'A'. But, as noted earlier, there is usually a good deal of overlap in conceptions for names that are very common, and names that have no referents can only be passed on by description, so they are especially

<sup>8</sup> There is a fourth argument (that is not classical) from the need to understand how intentional attitudes are to be described so as to play their usual role in psychological explanation. This need is discussed in my (2000: ch. 12).

likely to have conceptions that are largely public (though perhaps shifting).

Third, intensional contexts. A well-known way of extensionalizing intensional contexts was suggested by Davidson in 'On Saying That' (1968–9). His idea was that a sentence such as 'Galileo said that the earth moves' is true just in case uttering the words inside the 'that . . . ' clause of this sentence makes the speaker and Galileo into 'same sayers', people who have uttered words with the same import. I have adopted a similar view but generalized it, claiming that when one representation is held up or put on display in order to show what another representation is like, the kind of similarity intended may concern any aspect of meaning, or may even concern some aspect of the vehicle of the displayed expression (Millikan 1984: ch. 13; 2004: ch. 7). For an example of the latter, consider 'John kept insisting that it wasn't a woodchuck but rather a groundhog'. It is clear here that the similarity intended must concern the very words 'groundhog' and 'woodchuck', since these two are names for the same.

Intentional attitude contexts yield to a similar analysis. In 'John firmly believed that it was not a woodchuck but a groundhog' again something about the words 'woodchuck' and 'groundhog' is surely at stake. One way to understand this is to assume that an embedded sentence displayed in an intentional attitude context refers to an intentional attitude that is relevantly like one it would be its own stabilizing function to produce. The sentence 'It was not a woodchuck but a groundhog', if it were to serve its stabilizing function with John as a hearer, would produce just the mental state John is in, right down to the last conceptual component. That the message concerns not merely some proposition associated with John's mental state but also his conceptions, including the very words through which he would try to recognize information coming in about the subjects of his thought, is clear on the (pragmatic) assumption that John does not think a thought that shows, from the *inside*, that it is contradictory. This reading also nicely accommodates the fact that definite descriptions appearing inside intentional attitude contexts are sometimes read as attributing the description to the thinker as part of his conception and sometimes as attributing to him only a thought of the description's referent. Thus, 'Ralph thought that our venerable dean was a spy'

might or might not imply that Ralph knew that the man he thought was a spy is our dean. This is entirely natural if the intentional attitude description works by displaying a sentence whose function is to produce an attitude like the one being attributed and if definite descriptions have alternative stabilizing functions corresponding to Donnellan's distinction, as discussed above.<sup>10</sup>

The three aspects of meaning that I have discussed are thus sufficient to account for those properties of natural language traditionally associated with meaning.

## REFERENCES

- Boër, S. E., and Lycan, W. G. (1986), *Knowing Who* (Cambridge, Mass.: MIT Press).
- Davidson, D. (1968–9), 'On Saying That', *Synthese*, 19: 130–46; repr. in A. P. Martinich (ed.), *The Philosophy of Language*, 4th edn. (Oxford: Oxford University Press).
- Donnellan, K. (1966), 'Reference and Definite Description', *Philosophical Review*, 75: 281–304.
- Grice, H. P. (1968), 'Utterer's Meaning, Sentence-Meaning and Word-Meaning', in *Studies in the Way of Words* (Cambridge, Mass.: Harvard University Press).
- Lewis, D. (1969), *Convention: A Philosophical Study* (Oxford: Blackwell).
- (1975), 'Languages and Language', in K. Gunderson (ed.), *Language, Mind and Knowledge* (Minneapolis, Minn.: University of Minnesota Press); repr. in H. Geirsson and M. Lososky (eds.), *Readings in Language and Mind* (Oxford: Blackwell, 1996), 134–55.
- Millikan, R. G. (1984), *Language, Thought and Other Biological Categories* (Cambridge Mass.: MIT Press).
- Millikan, R. G. (2000), *On Clear and Confused Ideas* (Cambridge: Cambridge University Press).
- (2004), *Varieties of Meaning* (Cambridge, Mass.: MIT Press).

<sup>10</sup> To accord with this treatment, modal contexts need to be understood as contexts in which the truth conditions concern representations, rather than possible situations or worlds.

- Sellars, W. (1956), 'Empiricism and the Philosophy of Mind', in K. Gunderson (ed.), *Minnesota Studies in the Philosophy of Science*, i (Minneapolis, Minn.: University of Minnesota Press), 253–329; repr. in Sellars (1963).
- (1963), *Science, Perception and Reality* (New York: Humanities).
- Strawson, P. F. (1974), *Subject and Predicate in Logic and Grammar* (London: Methuen).

## 4 The Son and the Daughter: On Sellars, Brandom, and Millikan\*

---

Brandom and I were both Sellars's students, yet large differences have arisen between us. How have two siblings, both admirers of the father, come to stand so far apart? Which of us has abandoned the faith? Or was there a crack in Sellars's own position into which we have but driven a wedge?

It was with the latter in mind that I first approached this chapter. I thought that a crack might lie in the bridge that Sellars attempted to build between Wittgenstein's *Tractatus Logico-Philosophicus* and his *Philosophical Investigations*, both of which works Sellars admired greatly and which he claimed were not incompatible in basic measure. Certainly, there are 'Tractarian' themes in Sellars that only I have pursued and themes from the *Investigations* that only Brandom has pursued. I have pursued the picturing themes from the *Tractatus* that were carried through in Sellars's discussions of that causal-order relation between language and the world that he called 'representing' (e.g. Sellars 1963a; 1979: ch. 5). Also, in his unique interpretation of Kant, according to which the phenomenal world is abstractly isomorphic to the world in itself. Brandom has followed Sellars's interest in the language-games metaphor from *Philosophical Investigations*, expressed in Sellars as a form of inferential role semantics and in the thesis that one learns to think only as one learns to abide by the rules of a language. But on inspection there is, at least, no obvious crack in the bridge Sellars built between the *Tractatus* and the *Investigations*.

\* This essay also appears in *Pragmatics and Cognition*, 13.1 (2005).



Indeed, Sellars went to great pains to explain exactly how inferential role semantics was consistent with 'Tractarian' picturing. The idea was, roughly, that in an individual's or a community's following the rules of a language, the language being largely internalized as thought, a very abstract map of the world was in the process of construction.

A language, in its primary mode of being, simply is the pattern of beliefs, inferences and intentions (1979: 129)

in my account, the *manner* in which the names 'occur' in the 'picture' is not a conventional symbol for the *manner* in which the objects occur in the world, limited only by the abstract condition that the picture of an *n*-adic fact be itself an *n*-adic fact. Rather, as I see it, the manner in which the names occur in the picture is a projection, in accordance with a fantastically complex system of rules of projection, of the manner in which the objects occur in the world. (1963a: 215; 1979: 139)

These fantastic complexities are introduced mainly by the inference rules, formal and, more importantly, material, that govern 'statement–statement' (hence judgment–judgment) transitions. Just as the generalizations in question do not, so to speak, *separately* relate 'red' to red things and 'man' to men [but] relate *sentential* expressions containing 'red' to red things and *sentential* expressions containing 'man' to men (1979: 70)

so

[T]he representational features of an empirical language require the presence in the language of a [whole] schematic world story (1979: 128).

The map of the world produced by a language is not found sentence by sentence but only in the whole of the living *language cum thought* running isomorphically to the whole world in sketch. If there is a crack in the Sellarsian foundation, this is not where it lies, or anyway not precisely.

Where there may be a crack, however, is in Sellars's treatment of the nature of linguistic rules and the relation of these to conceptual roles and thus to intentionality. Conceptual roles for Sellars (as for Quine) were internalized patterns of linguistic response, responses to the world with words, responses to words with more words, and responses to words with overt actions. These patterns were not merely

patterns in fact, however, patterns actually engaged in by thinkers, speakers, and hearers. Sellars took linguistic rules to be normative rather than merely descriptive of regularities. Moreover, they were normative in a very strong prescriptive or evaluative sense. He was fond of saying that these rules were 'fraught with ought'. They prescribed regularities rather than merely describing them. He also compared these rules to the rules of a game (such as chess) in which conventionally allowable moves are made, the outcomes of which get counted, in accordance with further conventions, as having certain results. (Moving your rook to make that kind of configuration counts as putting my king in check.) His understanding of linguistic rules thus made contact with theories of speech acts that take these to be wholly conventional in the sense that acts of this sort could not be performed at all were there no conventions for performing them. The relevant norms are essentially social in origin and function.

On these various points about linguistic rules I think that Brandom pretty much agrees. True, he prefers to speak of 'practices' rather than 'conventions', but that is because others have analyzed the notion 'convention' as though all conventions rested on complexes of prior beliefs, reasons, and intentions, and Brandom, like Sellars, holds that beliefs, reasons, and intentions are themselves made possible only as a result of the relevant practices (1994: 232–3).

An obvious question concerns the relation of normative rules governing language and thought to actual regularities (hence to the actual picturing) found in language and thought. Sellars held that the linguistic rules were inculcated in children by socialization, which he took to be achieved by conditioning. Conditioning at first produces merely 'pattern governed behavior', but ultimately, through the introduction of metalinguistic patterns, also 'rule obeying behavior'.

To learn pattern governed behavior is to become conditioned to arrange perceptible elements into patterns and to form these, in turn, into more complex patterns and sequences of patterns. Presumably, such learning is capable of explanation in S–R [stimulus–response] reinforcement terms, the organism coming to respond to patterns as wholes through being (among other things) rewarded when it completes gappy instances of these patterns. Pattern governed behavior of the kind we should call 'linguistic' involves

'positions' and 'moves' of the sort that *would be* specified by 'formation' and 'transformation' rules in its metagame if it *were* rule obeying behavior (1963a: 327)

It is not in the first instance rule-obeying behavior, however, because the young child does not yet conceptualize the rules that it follows. Conceptualizing the rules is achieved by coming to use metalinguistic normative forms:

what we need is a distinction between 'pattern governed' and 'rule obeying' behavior, the latter being a more complex phenomenon which involves, but is not to be identified with, the former. Rule obeying behavior contains, in some sense, both a game and a metagame, the latter being the game in which belong the rules obeyed in playing the former game as a piece of rule obeying behavior. (ibid.)

Learning the use of normative expressions involves . . . acquiring the tendency to make the transition from 'I ought now to do *A*' to the doing of *A* . . . it could not be true of a word that 'it means *ought*' unless this word had motivating force in the language to which it belongs. (ibid. 350)

Thus, normative rules, for Sellars, are not translatable into nonnormative terms. Accepting a normative rule is not believing a fact but tending to be motivated in a certain way. Similarly, the work that the thought 'Now fetch some water!' does within one's psychology is not the work of a belief but the work of tending to cause one to fetch some water.

On the other hand, although to accept a norm is not just to know a fact, on Sellars's view the presence of normative rules in the natural world appears in the end as just one more level of fact in that world. From the scientific realist's standpoint, you can understand the nature of the normative practices of a community without participating in them. Similarly, although Sellars insists that the use of semantic and logical terms engages one in certain community practices, in semantic assessment, and so forth, it is also possible to understand what the functions of semantic and logical terms and statements are apart from being oneself engaged. You can understand these statements without participating in the practice of semantic assessment. It is one thing to *use* semantic language, for example, to say and mean or to understand ' "rot" means red'. But you can also *describe* the use of semantic

language without using it. You can describe what patterns of response in a language community, along with the origins of these responses in a history of language training, and training of the language trainers, and so forth, constitutes that 'rot' means red in that community. You can understand what the 'means' rubric does without indulging in it. You can understand specific forms of semantic assessment without participating in the particular practices being examined. There are *truth-conditions* for ' "rot" means red' of a perfectly ordinary, if very complicated sort. It's just that it's not the job of the sentence ' "rot" means red' to impart the information that these truth-conditions hold. Rather, its job is to get one to use 'rot' as one already knows to use 'red'.

Putting things bluntly, it seems that Sellars understands accepting semantic norms as merely displaying certain dispositions, dispositions to make certain moves in language and thought and dispositions to sanction these moves in others. Brandom claims that this sort of analysis will not do:

There clearly are socially instituted norms of this sort. Whatever the Kwakiutl treat as an appropriate greeting gesture for their tribe, or as a correctly constructed ceremonial hut, is one; it makes no sense to suppose that they could be collectively wrong about this sort of thing. The question is whether conceptual norms ought to be understood as of this type. (1994: 53)

In the case of conceptual norms,

assessing, sanctioning, is itself something that can be done correctly or incorrectly. (ibid. 36)

Defining normative attitudes in terms of dispositions to apply sanctions does not by itself reduce the normative to the nonnormative—it just trades off one sort of norm for another. (ibid. 42)

contents conferred on sentences by the score keeping practices I describe are not equivalent to the contents of any claims about what *anyone takes* to be true . . . their truth does not reduce to what I, or anyone else, or even *everyone* is or would be disposed to claim (1997: 202)

a cardinal criterion of adequacy of any account of the *conceptual* norms implicit in discursive practice is that it make intelligible their objectivity (1994: 63)

On a dispositional account of semantic norms we could not make sense of the fact, for example, that ‘it could be true that the sun will collapse whether or not everyone always thinks that it won’t’.

Brandom concludes that there must be ‘disposition transcendent conceptual norms’ and takes on as his central project to explain how this can be so. The alternative he sees to reducing the normative to dispositional terms is to posit that it’s ‘norms all the way down’. In setting out this position he remains committed to the Sellarsian view that the rules of language and norms of thought are instituted in public ‘practice’.

only communities, not individuals, can be interpreted as having original intentionality . . . the practices that institute the sort of normative status characteristic of intentional states must be *social* practices (ibid. 61)

Now, I agree with Brandom that conceptual norms must be disposition-transcendent, hence with his rejection of Sellars’s view of norms as derived from meta-dispositions to sanction. One wonders, however, whether the game metaphor with talk of ‘score keeping’ is really worth preserving after this insight, why it will not just prove misleading. Surely if everyone counts a certain move in chess as producing checkmate, or a certain move in basketball as scoring two points, ‘it makes no sense to suppose that they could be collectively wrong’ about these matters. There must be a deep divide between language and ordinary games that we should try not to obscure with a metaphor but instead to keep in full view.

Indeed, there is a competing theme in Sellars’s discussion of linguistic rules that seems both to introduce these rules as disposition-transcendent and to compete with the idea that these rules are at all like the rules of a game. He characterizes his position in part as follows:

when you describe the process whereby we come to adopt the language of which [some inferential] move is a part, you give an *anthropological*, a (very schematic) *causal* account of how language came to be used . . . in which you stress evolutionary analogies and cite the language of the beehive (1963a: 353).

In the case of bees,

(a) The pattern (dance) is first exemplified by particular bees in a way which is *not* appropriately described by saying that the successive acts by which the pattern is realized occur *because of the pattern*.

(b) Having a 'wiring diagram' which expresses itself in this pattern has survival value.

(c) Through the mechanisms of heredity and natural selection it comes about that all bees have this 'wiring diagram.' (ibid. 326)

In the case of humans,

the phenomena of learning present interesting analogies to the evolution of species . . . with new behavioral tendencies playing the role of mutations, and the 'law of effect' the role of natural selection (ibid. 327)

The analogy with bee dances retains the theme that conforming to the rules of a language is an intrinsically social activity. A bee dance is of use only if sister bees watch it and follow its direction. But the implication is clear that coming to follow the patterns prescribed by the rules of one's language community is not *just* a game but has some broader utility for the child or for its community. It has a value beyond that of displaying certain social graces (say, as in playing a decent game of chess or bridge in some social circles). Moreover, it is hard to believe that Sellars has overlooked that a bee dance is a tiny map of the location of some nectar. The bee dance not only has utility for the bees, but the fact that it maps the location of nectar by a certain rule of projection helps to explain *why* or *how* it can have this utility. It helps to explain the mechanism involved.

It is this second and, I believe, opposing metaphor of Sellars's that I have adopted in my work. The norms for language are uses that have had 'survival value', as Sellars put it. As such these norms are indeed disposition-transcendent, but they are not 'fraught with ought'. They are not prescriptive or evaluative norms. Their status has nothing to do with anyone's assessments. A norm is merely a measure from which actual facts can depart; it need not be an evaluative measure. A mere average, after all, is also a kind of norm. Behavioral forms that have had past survival value are a measure from which actual behavioral dispositions, both past and present, can depart, but such departures are in no sense proscribed. Indeed, departures sometimes

prove advantageous. What a biological or psychological or social form has been selected for doing, through natural selection, through learning, or through selection for social transmission, is a norm against which the form's actual performances can be measured. It is the 'natural purpose' of the form to fulfill this function, purposes, like norms, being *essentially* things that are not always fulfilled:

Contingencies may block the road of inquiry, yet truth (adequacy of representation [mapping] ) abides as the *would be* of linguistic representation. (Sellars 1979: 130)

The possibility of departure is built into the very notion of a *would be* or purpose. But to say that a natural purpose has not been fulfilled is to proffer an ostensible fact, not an assessment.

Looking carefully for the natural purposes of language and thought, however, reveals that these purposes cannot be all on one level. Both Sellars and Brandom see language and thought as a seamless whole. For Sellars, thought is just as inseparable from its expression in language as language is from the thought it expresses; the functional roles of language and thought each extend to include the other. For Brandom, the objectivity of conceptual norms derives from public linguistic practice. The original impulse for this idea comes from *Philosophical Investigations*, in the claim that the criterion for having followed a rule can only be public agreement. And surely something analogous to public agreement is required to keep the bees dancing. But if we ask whether the survival values of the concepts we acquire from learning a language are at root benefits gained only through the community by means of social cooperation, the answer seems to be no. Clearly there are benefits to the isolated individual as well. Conforming to the semantic rules embodied in a language is not just a social activity, of use only within a society. If learning a language is learning to think, having learned a language will also come in handy on Robinson Crusoe's island, with or without assistance from Man Friday. Playing a conceptual game of solitaire must also have its advantages. But then there must exist standards of conceptual clarity accessible within individuals apart from the language community, standards by which merely wrestling with

nature determines when a useful conceptual pattern has been formed. Whether one's thought is well formed has a criterion that also applies when one is alone in one's workshop. The bee that dances correctly can follow its own dance to nectar.

It need not follow that the functions of language derive from functions that thought intends for language. We need not follow Stalnaker in 'dividing up the fundamental orientations of various approaches to intentionality, accordingly as rational agency or linguistic capacity is taken as primary' (Brandom 1994: 149). If language and thought do not form a seamless whole, that doesn't have to mean that either Grice is right and the intentionality of language derives from that of thought, or that Wittgenstein is right and the intentionality of thought derives from that of language. Selection takes place on various levels. Most obviously there is selection of genes, selection of behaviors by conditioning and by trial and error learning, and selection of traits and behaviors for social transmission. Each of these levels produces its own yield of natural purposes. The selection of language forms takes place on the social level. Language survives when it serves cooperative functions often enough, functions that reward at once both speakers and hearers (though they may often be rewarded at the end in different ways). Language forms proliferate when aiding speaker and hearer cooperation on common projects—typically, the sharing of information speaker and hearer have a mutual interest in sharing or the coordinating of projects and activities they have a mutual interest in advancing. Language is something that it takes a pair of people to do; both must be purposefully involved. Completed speech acts of a kind that have survival value are not the work of a speaker alone, but of a hearer purposefully cooperating with a speaker. Purposeful doings need not be confused with doings guided by intentions, however. There is purpose in what the kidneys do and purpose in the exhibition of behaviors resulting from conditioning. That producing beliefs or desires in a hearer is often part of the natural purpose of language use, both a purpose of the speaker's speaking and a purpose of the hearer's reaction in understanding, does not require that either speaker or hearer has intentions concerning beliefs or desires or, indeed, so much as concepts of beliefs and desires (Millikan



1984: ch. 3; 2004: ch. 9). Surely Sellars was right that speaking comes before thought about thought (1963*b*).

To say that the use of language results in acts of a special kind that it would be impossible in principle to perform outside the conventions or the practice of language use is misleading, however. In its cooperative way, language accomplishes perfectly natural things. For example, doing something that produces a certain belief or intention in another is a perfectly ordinary thing to do. That it can be done cooperatively through the use of language does not change this matter. Conventions, on this view, are merely ways of doing things that are proliferated by being reproduced, and that exhibit a certain arbitrariness of form. They are reproduced patterns that proliferate due partly to weight of precedent, rather than due, for example, to intrinsic superiority (Ch. 1). The conventions of language do not create any new kinds of action effects. Language conventions are best thought of merely as lineages of behavioral patterns involving a speaker's utterance and a hearer's response. They do not correspond to rules, and certainly not to prescriptive rules. It is true that many conventions are ways of doing things to which one ought to conform, given that there are such conventions. For example, conventions about which side to drive on and whether to stop at the red or the green are conventions with which one ought to conform. Moreover, in traditional cultures, doing things in unconventional ways is often proscribed quite generally. But this evaluative kind of normativity is something added to mere conventionality. Decorating for Christmas with red and green is conventional, but surely in no way required. Conforming to the conventions, engaging in the linguistic practices of the community in which one lives, is, in the main, merely a practical matter. Mainly it concerns how to accomplish certain practical tasks in a given environment (Ch. 1).

But, as Brandom has said, 'a cardinal criterion of adequacy of any account of the conceptual norms . . . is that it make intelligible their objectivity'. How do we do that without appealing to linguistic practice? By what objective criterion can one be following a rule of thought privately, following in a way that no one else will assess or, indeed, even notice or care about? What objective criterion determines that one is

using a dog thought only in response to dogs or that one's dog thoughts always correspond even to the same kind of thing?<sup>1</sup>

I adopt Sellars's suggestion that adequate intentional representing is a kind of picturing or mapping. And I adopt his suggestion that this picturing or mapping may have immediate practical uses, as when one bee makes a dance map that guides another toward nectar. The suggestion is then that the functions of both bees might be realized within the same brainwork, one part of the network making maps of the world that will guide the other in directing behaviors for navigating that world. This first and simplest model for cognition gives us perception directly for action, perception-action cycles, roughly as conceived by Gibsonians. Representations or 'icons' that directly mediate between perception and action I call 'pushmi-pullyu' representations. Like the bee dance, they at the same time tell what the case is with some part of the world and direct what to do about it. Behaviors of the very simplest animals are governed by pushmi-pullyus of this kind, as are myriad automatic responses of humans to the most immediate environmental contingencies facing them, such as being off balance or needing to navigate rough or smooth ground or needing guidance by perception in performing routine motions, for example in grasping and manipulating objects.

Already at this simple level a stringent criterion of correctness for rule-following is in effect. The perceptual systems must manage systematically to deliver representations of the world that accord with a rule of correspondence to which the action systems are also adjusted. On the Wittgensteinian picture, one language user trains another, the evidence that each is conforming to a rule being that their results match. Similarly, that both the perceptual systems and the action systems are conforming to rules is evidenced by the fact that the results of their cooperative activities on varying occasions are constant. The bees get to nectar, the body remains upright, the path on the ground is negotiated, and the coffee cup safely transported to the mouth. Moreover, consistent conformity to rules at this level is a very

<sup>1</sup> The sketch given below is developed in my (1984, 1993, 2004), coupled with the epistemology of theoretical concepts developed (most fully) in my (2000).

considerable achievement. The perceptual systems must locate the layout of distal circumstances through a wide variety of mediating conditions, such as frequently changing lighting conditions, visual static, occluding objects, changes in position of the body and eyes, and so forth. They must recognize the same individual or the same kind or the same stuff again, so as to represent it consistently, from numerous angles, perhaps as in numerous postures, manifesting itself in a variety of ways through different sensory modalities (2000, 2004: pt. IV). To make perceptual maps for action that map consistently, recognizing relevant perceptual constancies, showing forms and objects by rule in a consistent way, is a task of enormous complexity.

Beyond perception for action, humans, at least, make cognitive maps that are not dedicated in advance to the guidance of particular behaviors. We collect great quantities of information with no immediate uses in view, storing it away for possible later contingencies. Having separated the descriptive from the directive aspects of representation, these have to be joined together again through practical inference. But representations of fact that are not immediately tested in action and that are then used to form more representations and then still more through inference need to be screened for accuracy and consistency in some way. Rules or patterns of belief-formation need to be strictly regimented as they are developed, well in advance of practical uses for the resulting beliefs. Wittgenstein proposed that this screening is accomplished by the criterion of agreement in judgments with others. I have proposed that it is done in primary instances by the criterion of agreement with one's self in judgments. Agreeing judgments need not be made by different persons. Judgments can be made by the same person in different ways, from different perspectives, under changing conditions, using different sensory modalities, employing different inferential patterns (1984: ch. 18–19; 2000: ch. 7; 2004: ch. 19). Agreement with oneself in judgment attests to the fact that one is managing to map again the very same objective structures in the world through different methods of projection. Indeed, agreement with others is discovered only as a form of agreement with oneself. Agreeing with others is not speaking in unison. If you and I say in unison 'That cookie is mine', we are disagreeing. To recognize agreement with another in

judgment, you have to advance for yourself rules of translation by which another's speech carries information to you, these rules being entirely parallel to the patterns or rules in accordance with which you translate sensory information arriving through a great variety of other media into beliefs (2000: ch. 6; 2004: ch. 9).

This being said, nonetheless there remains something very special about agreement with others in judgments. We acquire the vast majority of our concepts through the medium of public language, just as we acquire the vast majority of our practical and social skills from others. But, far more important, the larger proportion of our concepts could not in principle have been developed solo because the multiple perspectives and sources of information required to test their objectivity are made possible only through cooperation with others who have independent access to the same objective affairs through other temporal and spatial perspectives. To take just one instance: concepts of dated occurrences, indeed all concepts involving historical time, would seem to be possible only with the help of others informed of these occurrences independently from other perspectives (2004: ch. 19).

That is quite enough said about discord between the son and the daughter. Let me end by noting a very deep theme that is common to Brandom's work and mine, binding us together and setting us apart from others currently writing about language and thought.

Brandom and I are both committed to explaining the meanings of linguistic expressions in terms of their use (though there are, of course, differences in how we understand 'use'). Brandom puts this by saying that 'semantics must answer to pragmatics' (1994: 83), arguing against 'representationalism': the view that representation comes first, then inference, then use. A representation is something that purports to represent, and purporting to represent is purporting to represent to some interpreter, some user who is, 'taking, treating, or using a representation as a representing' (ibid. 75). MacDowell puts the position this way: 'We cannot work up from the semantics of words to the semantics of sentences, and only then move up to consider the structure of the language game' (1997: 158).

I have taken exactly the same position, though the terminology is different. In the case of language, what Brandom calls the study of

'pragmatics' corresponds roughly to what I call the study of 'function'. 'Meaning', in the most basic sense, simply *is* function; it is what I have called 'proper' or 'stabilizing' function (1984: chs. 1–6; 2004: chs. 2, 11), or, very roughly, what Sellars called 'survival value'. The functions of complete linguistic forms are to perform complete speech acts, these being cooperative acts accomplished by speaker and hearer together. The performance of cooperative acts is what keeps speakers using these forms in consistent ways and hearers responding to them in consistent ways—hence keeps them in circulation. I have tended to reserve the term 'pragmatics' for the study only of how nonconventional speech acts are performed, acts which do not express conventional functions because not directly derived from precedent, but this is a mere terminological difference. The 'semantic' dimension of representation, if we understand by this the involvement of truth- or satisfaction-conditions, is owed to a certain *way* of performing a function, a certain kind of mechanism that is employed. Satisfaction-conditions are related to function as a method or manner is to a performance.

In the case of thought, I have argued, there is no intentionality prior to the emergence of complete representations having truth- or satisfaction-conditions, and representations cannot have satisfaction-conditions unless they have uses. Briefly put, there is no such thing as intentionality without attitude. Participating in inference processes by which new descriptive and directive representations are formed is a central way in which human intentional attitudes are employed, so the intentionality of these attitudes and their content is a function, in part, of inferential patterns. That these patterns have to match their content and that their content depends in part on these patterns are two sides of a coin. On the other hand, I have claimed, there is also plenty of intentionality prior to that of the intentional attitudes, both in the perceptions of animals and humans and in simpler messenger systems that abound in the body.

## REFERENCES

- Brandom, R. (1994), *Making it Explicit* (Cambridge, Mass.: Harvard University Press).

- (1997), 'Replies', *Philosophy and Phenomenological Research*, 57/1: 189–204.
- McDowell, J. (1997), 'Brandom on Representation and Inference', *Philosophy and Phenomenological Research*, 57/1: 157–62.
- Millikan, R. G. (1984), *Language, Thought and Other Biological Categories* (Cambridge, Mass.: MIT Press).
- (1993), *White Queen Psychology and Other Essays for Alice* (Cambridge, Mass.: MIT Press).
- (1998), 'Language Conventions Made Simple', *Journal of Philosophy*, 95/4 (Apr.), 161–80 (repr. as Ch. 1 above).
- (2000), *On Clear and Confused Ideas* (Cambridge: Cambridge University Press).
- (2004), *Varieties of Meaning* (Cambridge, Mass.: MIT Press).
- Sellars, W. (1963a), *Science, Perception and Reality* (New York: Humanities).
- (1963b), 'Empiricism and the Philosophy of Mind', in *Science, Perception and Reality*.
- (1963c), 'Truth and Correspondence', in *Science, Perception and Reality*.
- (1979), *Naturalism and Ontology* (Atascadero, Calif.: Ridgeview).

## 5 The Language–Thought Partnership: A Bird’s Eye View\*

---

In this chapter I will try to sketch in miniature the whole of my work on the relations between language and thought. I have offered close-ups of this terrain in various papers and books, to which I will refer freely. Here I will focus on the relations among the parts, hoping to provide a short introduction to my work on language and thought for some, and for others a clarification of the larger plan.

I take language and thought to stand largely parallel to one another. For example, the intentionality of each is defined independently of that of the other: thought is possible without language, and language is possible that does not convey thought. On the other hand, public language is not merely a stimulus to the development of thought. It is constitutive of developed human thought. These are the relations I hope to clarify.

The first idea needed to understand the relation of language to thought, I claim, is that of a ‘proper function’. Linguistic meanings are (in the first instance) the proper functions of various language forms. Further, the intentionality both of language and of thought concerns the way these serve their proper functions. Speaking extremely roughly at first, proper function is survival value. In the simplest instances, the proper function of an item is what it has been doing—better, what the predecessors from which it was reproduced or copied were doing—that helped account for continued reproduction or copying and hence for its existence. In less simple instances the proper function must be described as varying in relation to obtaining

\* First printed in *Language and Communication*, 21 (2001), 157–66.

conditions, or the function is derived from the functions of prior devices that function properly by varying their outputs according to conditions. I defined 'proper function' and discussed implications of this definition in Millikan (1984: chs. 1, 2; 1993: chs. 1, 2).

When speaking informally I used to call proper functions 'biological functions', but this led to misunderstandings. It was thought that my analysis was an attempt to analyze biologists' usage of the term 'function', and it was thought that I held that all proper functions are derived from genic selection. Neither of these is right. Especially obvious, the functions of the elements of specific languages, such as words, syntactic forms, tonal inflections, and so forth, are not filtered through the genes. More generally, the functions of learned behaviors are related only in indirect ways to the genes that account for the dispositions to learn these behaviors under specific conditions, and the functions of thoughts are related very indirectly to the genes that account for the abilities to develop the concepts exercised in them. These issues are treated in Millikan (1984, 1993, 1995, 1996, 2000*b*) and I'll say no more about them here.

Intentionality has to do with truth-conditions or, more broadly, satisfaction-conditions, which I will discuss in time. At the moment, I want only to clear the ground by saying what intentionality is not. It does not correspond to a peculiar kind of proper function that intentional items like sentences and beliefs can have (e.g. the function of 'indicating' or 'representing' or the like). Rather, it corresponds to a special *manner* in which some things serve their proper functions. I will explain this later. Right now I want to put intentionality aside and talk just about the functions of language forms.

Not all complete functional units of public languages have satisfaction-conditions. Some just have proper functions. 'Hello!' has no satisfaction-conditions, but it has a function. What, exactly, is involved in a language form having a function? Its proper function is its survival value. It is something that it does or has been doing that accounts for the fact that it continues to be repeated, reproduced, so that it does not die out of the language. Usually, what language forms do in this capacity is to produce a characteristic response in a hearer or hearers. (Elementary language forms do not do this by themselves, of



course, but in combination with other language elements functioning together in ways admitting of systematic description.) Why does the language form continue to produce this characteristic response? We must ask the same question about the response produced as about the form itself. What is its survival value? Why does it keep recurring? Why does its lineage not die out?

I call a proper function of a meaningful form in a public language a 'stabilizing function'. This is a function that, when performed, tends both to encourage speakers to keep using the device and hearers to keep responding to it with the same (with a stable) response. A stabilizing function of a language device must be one that accords both with the speaker's purposes and with the hearer's purposes often enough to keep the device in circulation along with a stable response to it. Stabilizing functions are found where speakers and hearers cooperate, where their immediate interests overlap, though overlapping immediate interests sometimes conflict, of course, in ulterior motives.

The easiest illustrations of stabilizing functions are the functions of syntactic forms constituting the indicative, imperative, and interrogative moods. Roughly, indicatives have the function of producing or activating true beliefs in the hearer, imperatives have the function of producing hearer compliance, interrogatives have the function of eliciting true answers (details are in Millikan 1984: ch. 3). Stabilizing functions for a variety of other kinds of language devices are discussed in Millikan (1984: pts. I, II; Ch. 8; Ch. 9). The mechanisms at work in stabilizing the functions of public language devices are, in certain crucial respects, much like those at work in biological evolution under natural selection. Especially clear is the similarity between stabilization of the functions of various public-language forms through social selection processes and stabilization of the functions of animal signs, such as mating displays, danger signals, territory markers, bee dances, and so forth, through genetic selection processes. Notice that in both cases the actual *physical forms* that survive, as opposed to the functions they perform, have a high degree of arbitrariness.

Biological traits are often selected for functions they perform not on the average but just often enough to be worth preserving them for. For example, the eye-blink reflex may well be triggered more often by

other stimuli than by objects that would actually have hit the eye. Similarly, stabilizing functions of language forms are not necessarily statistically average functions or even common functions. A language device needs to perform its stabilizing function only in some critical mass of cases, the proportion varying with the average positive value of the function over against the costs of function failure for speakers and for hearers. We all know the effect even of very occasional weakness when dealing with a child who teases for things. The disposition to tease can survive many failures. Similarly, it is no mystery that the forms and idioms of daily speech can withstand much misuse, many failures of cooperation, and use in a multitude of parasitic or secondary ways, without altering their stabilizing functions.

Another similarity between stabilizing functions of language devices and biological functions is that neither is defined by reference either to anyone's intentions for them or anyone's cognition of them as having functions. Blinking one's eyes in response to an approaching object has a purpose, a proper function, but this function corresponds to no one's intention. Using smiles as a reward, it is possible to condition a person to blink more frequently without their awareness that they are doing so or of the purpose the blinks are serving. Their learned blinking behavior has a proper function of which they are unaware. Where actions are backed by explicit intentions, moreover, they often have intermediate purposes that are not explicitly intended. When I turn the key in the ignition of my car I intend to start it, but an intermediate proper function of my learned behavior is to cause electricity to flow in certain wires and a magnetic field to materialize in the starter motor, and so forth, all of which I may know nothing about. Similarly, one function of a child's use of 'More juice!' may be to produce the belief in her hearer that she desires to have more juice, yet the child may not possess so much as the concepts of belief and desire, let alone an intention to produce a belief about a desire. The notion 'stabilizing function' for language devices is defined by reference to purposes or proper functions of speaker utterances and of hearer responses, not in terms of speaker intentions and hearer understandings. (The mistaken notion that Gricean intentions underlie language use is discussed in Millikan 1984: ch. 3).

Thus, the definition of function for public-language devices, and hence, as remarked earlier, the primary definition of linguistic 'meaning', makes no reference to human intentions or, indeed, thought. It does not follow that the stabilizing functions of specific language parts can all be described without reference to thoughts. For example, I have described the specific function of the indicative mood as the production or activation of true hearer beliefs. Similarly, if it is the function of the imperative mood to produce compliant hearer acts, presumably it is also its function, prior to that, to produce consonant hearer intentions. On the other hand, not all complete and functional language forms have as their stabilizing functions to express, transmit, or impart intentional attitudes. Nor do I refer merely to greetings, exclamations, expletives, and so forth. There are complete sentences that do not function to convey intentional attitudes. But to explain what I have in mind here I must first discuss intentionality.

The question whether a language form displays intentionality is independent of the question whether its function involves transmission of intentional attitudes from speaker to hearer. Let me begin with an analogy from ethology, one I have often relied on because it can be used to illustrate so many important points at once. Consider the dance of the honey bee. This dance has a whole series of proper functions. It stimulates the nervous systems of watching worker bees in certain ways, causing them to fly off in a certain direction for a certain distance determined as a function of the shape and angle of the dance. This brings the worker bees to the location of flowers bearing nectar, producing stimuli which cause the nectar to be collected by the workers and brought back to the hive, providing nourishment for the next generation of bees. To initiate this whole series of effects is also the proper function (series of proper functions) of the mechanism that is responsible for producing the dance in the dancing bee. The mechanisms that produce bee dances and the mechanisms that respond to them are cooperating devices that have coevolved, resulting in stabilization of the dance form, its method of production, and its effect on watching bees. The analogy I have in mind, of course, is with human language forms that have become stabilized in their cooperative functions. Now, a natural thing to say about bee dances is that they represent

the location of nectar which the dancing bee has discovered. What is it about them that makes it natural to call them representations of nectar locations?

First, notice that it is not their proper functions that make them representations. To have the function of initiating neural responses, of producing flight in a certain direction, of causing arrival at a site of nectar, of causing larval bees to be fed, and so forth, these functions, just as such, have no more to do with representing than do the functions of circulating blood or digesting food. Nor are they representations because their function is to convey intentional attitudes. Just to keep the biology honest, there is evidence that bees carry neural maps in their heads, but we do not have to assume this is so to see that the dances are representational. It could be that the watching bees responded directly to the dances by pivoting about to a certain direction and flying that way for a certain time.

What makes the dances into representations is not what they do but why they work, *why* they help to cause arrival at sites of nectar, hence arrival of nectar in the hive, hence well-fed larval bees. They work by bearing a correspondence to what they represent, according to a certain projection rule, such that in being guided by the dance the worker bees are caused to fly toward nectar. The 'certain rule of projection' describes the relation the dance must have to the location of nectar if it is to cause the watching bees to fly towards nectar when they react to the dance as they are designed to. The intentionality lies not in the function of the dance, but in the explanation of how the function is performed, in the principle involved. Roughly, the principle is mathematical isomorphism. Variations in possible bee dances to which worker bees are designed to respond correspond one-to-one to variations in possible locations of nectar in such a way that being guided by the dances produces arrival at sites of nectar.

Why is this intentionality? Because the dances display the characteristic trait of the intentional; namely, they can be wrong or false. They can fail to correspond as they should to a place where there is nectar. Should anything disturb the normal mapping between the shape of the dance and the location of nectar, this misalignment will, quite literally, lead the workers astray. Bee dances have truth-conditions.

The rules by which they are designed to correspond to nectar locations are semantic rules.

The intentionality of language is exactly parallel to the intentionality of bee dances. Language forms have, first, a function or series of functions. Next, we can ask how these functions are performed, what principles are involved. If, given the normal stabilizing hearer reaction to the form, the form will guide the hearer so that its stabilizing function is performed *only when there is a correspondence by a given rule or function between form and some structure in the world*, then the form is intentional. It has a truth-condition.<sup>1</sup> (For a full account of this see Millikan (1984: chs. 6 ff; 2004: chs. 5–6).)

For example, there are many indicative sentences of English the stabilizing functions of which are to produce true hearer beliefs, but that will not produce this result in a hearer who is responding in the stabilizing English-comprehending way unless they correspond to the world by certain mapping rules; namely, certain semantic rules of English. In this particular case, of course, the stabilizing function involves the production of intentional attitudes. But the general notion of intentionality for language has not been defined such that intentional language forms have to express or transmit intentional attitudes. And, indeed, in Millikan (1984) I argued that traditional puzzles about the meanings of a number of language forms, including sentences expressing identities, sentences explicitly asserting existence, and sentences of the form ‘*x* means *y*’, find solution when we allow that the sentences can be intentional, can have truth-conditions, without having as their stabilizing functions to produce intentional attitudes. I suspect this is also true, for example, of sentences expressing various of the modalities.

The kind of intentionality I have been discussing is the kind had by representations of fact. I call these ‘indicative representations’ or

<sup>1</sup> Typically, this will be true only of complete sentences. Portions and aspects of sentences that make a systematic contribution to truth-conditions can be considered to be intentional in a derivative way. That is, the intentionality of the complete representation which sports a truth-condition is prior to the intentionality of any of its parts or aspects. Truth-conditions are not built up from term references. Rather, term references are abstracted from truth-conditions.

'fact icons'. Representations of goals, 'imperative representations' or 'goal icons', are another matter. Consider the bee dance again. It is natural to say that the dance represents the location of nectar, but it is just as natural to say that the dance represents where the bees are to go, their goal. Again, the intentionality is there because of a correspondence rule or a mapping, but this time the rule *is* determined by a proper function of the dance. A proper function of the dance is to cause the watching bees to fly to a location that bears a certain relation to the dance; namely, the relation described by the correspondence rules for the dance, the semantic rules of B-mese. This is the paradigm for all goal representations, including directive sentences (Millikan 1984: esp. ch. 6; 1993: chs. 3–5; 2004: chs. 5–6). Again, as in the case of the bees, it need not be that the goal expressed by a directive sentence is normally reached by means of the interpreter's first forming a corresponding intention. That is not built into the definition, although in the case of explicitly directive language forms it does happen to be the case.

On the other hand, the bee dance is not really the best exemplar for either fact representations or goal representations. That is because it has two faces, whereas pure fact icons and pure goal icons each have only one. The bee dance represents the facts and tells what to do about them all in one breath. I have called representations of this kind 'pushmi-pullyu representations' (Ch. 9), arguing that they are found also in both human language and thought. Consider 'No, Johnny, we don't eat peas with our fingers' and 'This road is legally closed'.

The way the definition of 'proper function' is set up (Millikan 1984), it falls out that artifacts and human activities have as proper functions whatever functions their makers or performers intended them to serve (*ibid.*: ch. 2). Similarly, where nonrepresented proper functions or implicit purposes are involved in the making or performing, they have these purposes as their functions. Recall the function of the unconscious blinking that has been trained by smiles, and recall the various functions concerning parts under the hood that my turning of the ignition key has. Now, the maker (performer) of a public-language expression or sentence token is the speaker or writer of the token. So one proper function of a public-language token is whatever the speaker explicitly intended and/or nonexplicitly purposed that it

accomplish. But another proper function that it has is the stabilizing function of its type; roughly, its literal meaning. The former is a 'derived proper function', derived in this case from speaker's intention or purpose. The latter is a 'direct' and 'stabilizing proper function', which has been accounting for survival of the expression type (or its elements and their form of concatenation) in the language community. (For details of this distinction, see Millikan 1984: ch. 2). These two proper functions may be consonant with one another or, in more interesting cases, they may conflict.

A large portion of the subject matter of pragmatics concerns the interaction between speaker-derived functions and stabilizing functions of tokens of language forms. Stabilizing *uses* of any language form must be uses in which these two kinds of functions do not conflict; indeed, in which the derived function includes the direct stabilizing function as a part. The speaker purposes the literal function and the hearer cooperates (see Millikan 1984: esp. chs. 3, 4). Stabilizing functions concern the 'conventional' nature of language (Ch. 1, Ch. 8), and for the most part the conventional outcomes of speaker uses of language, the perlocutionary effects that conventionally follow these uses (in the mild sense that living together 'conventionally' follows marriage), accord with their stabilizing functions; that is, with their public-linguistic 'purposes'. Thus, in typical stabilizing uses of language, speaker purpose, public-linguistic purpose, and conventional outcome all coincide. But there are also many cases in which either one and two, two and three, or all three come apart. The modern debate about which if any speech acts are performed conventionally arises from this confusing overlap (coupled with inadequate theories about what language 'conventions' are: Ch. 1; Ch. 8).

To understand the intentionality of thought with care requires thorough understanding of the nature of 'derived proper functions', and I cannot say very much about that here. Still, the rough idea is not hard. The perceptual, cognitive, and conative systems of the human, we assume, have a normal way of developing from embryonic form to their mature adult state. This normal development is describable, however, only with constant reference to input from an environment that is 'normal' in a variety of respects—'normal', in particular, relative

to very broad features of the historical environment of the human species. Moreover, certainly in the case of the cognitive and conative systems, the description of the normal state of these systems at more advanced stages of development is a highly relational description. What is normal is for certain generally describable *kinds* of relations to have developed between these systems and certain *kinds* of structures in the environment, resulting in certain capacities and dispositions of the organism to interact with its specific environment, hence in behaviors adapted to that environment. But exactly which *particular* relations and capacities it is normal for a particular human to have developed depends on the particular input, the experience the system has had, the particular environment it has been adapted to. This much seems incontrovertible.

Now, how normal cognitive and conative development occur is, presumably, theoretically explainable. That is, there are principles, some perhaps quite specific, others very general, that account for the possibility of normal development, and account in a general way for the adaptiveness of the behaviors that result from it. Some of these principles must be very general principles indeed, such that they are operative over very wide variations in environmental input to produce what are, *relationally* described, the same *sorts* of adaptations of these systems and of the behaviors they control although to different environments. Were this not so, there could be no possibility of a science of normal developmental psychology or normal adult psychology. There would be no explanation for the fact that humans growing up in widely different environments tend to behave in ways adaptive in those environments. Put simply, I know my way home but you don't, nor would it help you much if you did. You know your own way home. It does not follow that there are not univocal general principles of learning that account for the development of both our abilities to get home, and univocal principles that explain how both are exercised, how our neural states function to get each of us home.

Now add another hypothesis. Many adaptive behaviors that are normal for organisms of various species are controlled by plastic states or structures within the organism that vary in a systematic way to parallel variations in the environment. The obvious simple cases here



are perceptual states of those organisms that can perceive aspects of their distal environments. Systematic parallel variation equals isomorphism of inner state to environment in accordance with some abstract rule. The perceptual state is, normally, aligned with the environment in accord with a definite alignment function. Suppose this to be true for cognitive states as well, for the states that correspond to intentional attitudes. One would expect the alignment or mapping functions to be much more abstract, of course. Fodor's vision of a 'language of thought' would be one among other ways to envision this very general possibility. All that would now be needed to demonstrate the intentionality (as intentionality was described earlier in this chapter) of intentional attitudes would be this assurance: reference to this sort of mapping helps to explain how the systems that manufacture and use intentional attitudes manage to perform their proper functions. That is, we have to assume that the cognitive systems of the adult are designed by evolution to perform in this sort of manner, or, better, that they were designed to *learn to perform* in this manner.

Now, it seems not impossible, at least, that when perception is used to guide bodily motion different members of the same species might use physically identical or similar perceptual representations to represent the same environmental features. They might use the same perceptual 'notation', as it were, or perceptual 'maps' having identical keys and projections. But this is not at all reasonable for human thought, because people can represent widely different things in thought, depending on their past experiences. They do not have ideas of the same things—not of the same people, nor the same kinds of objects and events, nor of the same properties. If thought involves mental representations, then each of us must think in his or her own individual representational scheme. There could not, for example, be such a thing as '*the* language of thought'.

A hugely interesting developmental question then concerns how the cognitive representational system of the individual human is developed. Ultimately, the question belongs to experimental psychology. The philosopher's job is to suggest directions in which to look, and how to recognize the desired object when sighted. I discuss this at some length in Millikan (1984: pts. III, IV; 2000) and will present

only a small fragment here, focusing on the way language enters the cognitive scene.

The problem that faces the cognizing organism is: (1) to develop a representational system with which to map relevant affairs in its world; and (2) to learn to make accurate representations of relevant aspects of its environment using this mapping system. The difficulty, of course, is that the aspects of the environment that it needs most to map are distal, and correspond in highly complex and often unreliable ways to proximal aspects. Correlatively, the organism has no direct access to what has to be mapped so as to tell when its maps are accurate. Speculations about how all this is done must begin with ontology: the notion of a map of a structureless world is incoherent. And they must proceed through epistemology—not the traditional epistemology of judgment but an epistemology of concepts. How do we know when our concepts are clear, when we are representing what in fact is the same *as* the same, being neither equivocal nor redundant as we form our basic representational systems? These projects I have tackled in Millikan (1984: pts. III, IV) and especially (2000) and I will not try to abbreviate here. But between ontology and epistemology lies the philosophy of mind, which I will talk about here. Granted a world with objective structure, what is it to have a concept corresponding to some aspect of that objective structure? What is it, for example, to have a concept of Fido or of dogs or of sugar or of round?

To have a concept of one of *those* sorts of things—to have, that is, a paradigm *empirical* concept—involves (caution: I do not say ‘equals’) a rough practical ability to reidentify that thing: ‘practical’ in two senses. The ability is ‘practical’ because it is fallible, and could not be otherwise even in theory. It is an ability that gets one by, for the most part, in practice. And it is ‘practical’ in that it must be applicable in practice. It must be a practical way of recognizing when the information contained in energies bombarding one’s senses is about that same thing again rather than about something different. This sort of ability typically includes numerous sub-abilities, corresponding to different ways of recognizing the concept’s object under a great variety of conditions, when bearing a great variety of relations to the thinker, and through a variety of intervening media. Nor are any of these sub-abilities more

basic, more 'definitional' of the concept, than any others (Millikan 1984: ch. 15; 2000). Surprisingly, by this route enters language.

Language enters as just one among the many other media by which information about the disposition, among other such things in the environment, of an empirical object, kind, stuff, or property can manifest itself to the senses. Just as there are conditions under which patterns of light striking the retina will vary systematically with certain properties of the distal objects reflecting it, there are conditions under which the sentences that a person hears will vary systematically according to the dispositions of things in the world that originated them. In neither case are the relations entirely dependable, and there is, of course, much more static in the case of sentences, the correspondence rules vary more radically when one moves from one surrounding language community to another than from one kind of lighting conditions to another. The basic principle, however, remains exactly the same.

The result is that a very large portion of our conceiving is done mainly or entirely through the medium of language (Millikan 2000: ch. 6; 2004: ch. 9). As individuals, each of us has empirical concepts that are entirely dependent on language, having no means of recognizing the objects of these concepts 'in the flesh'. We recognize them only through their manifestations in the speech of others. Even more important is the role of language in concept learning. Mastering the phonetic and phonemic structure of a language so as to be able to recognize, in general, when one has encountered the same word again is, quite literally, play for babies. It is much easier than learning how to reidentify 'in the flesh' the vast number of natural objects, kinds, and properties that concern us. Extreme examples are the many concepts from the sciences that were historically developed through great labor over long periods of time but are learned in minutes or hours when one studies the sciences. And with the stabilizing hand of language to hold on to, abilities to recognize the objects of these concepts in the flesh are often readily obtainable as well.

So which comes first, thought or language? The relations, if they are as I have outlined them, are very complex indeed.

## REFERENCES

- Millikan, R. G. (1984), *Language, Thought, and Other Biological Categories* (Cambridge, Mass.: MIT Press).
- (1993), *White Queen Psychology and Other Essays For Alice* (Cambridge, Mass.: MIT Press).
- (1995), 'A Bet with Peacocke', in Macdonald and Macdonald (eds.), *Philosophy of Psychology: Debates on Psychological Explanation* (Oxford: Oxford University Press), 285–92.
- (1996), 'Cognitive Luck: Externalism in an Evolutionary Frame', in P. Machamer and M. Carrier (eds.), *Philosophy and the Sciences of Mind* (Pittsburgh, Pa.: Pittsburgh University Press; Konstanz: Universitätsverlag Konstanz), 207–19.
- (2000), *On Clear and Confused Ideas: An Essay about Substance Concepts* (Cambridge: Cambridge University Press).
- (2004), *Varieties of Meaning* (Cambridge, Mass.: MIT Press).

## 6 Why (Most) Kinds are not Classes\*

---

In the last century many hundreds of experiments were run by psychologists trying to discover how people ‘classify’ or ‘categorize’ items under kind words such as ‘dog’, ‘chair’, and ‘fruit’. The position I have taken on such words is that they do not designate classes but units of another kind entirely. There do exist some, a very few, uncompounded nouns that designate classes, but words like ‘dog’ and ‘chair’ and ‘fruit’ are not among them. Here I will introduce you to this negative position. I cannot attempt to defend it at any length in a short chapter, but I will present a portion of it in the most intuitively understandable terms I can muster. The details are spelled out in *On Clear and Confused Ideas* (Millikan 2000), which I will refer to as ‘OCCI’.

One place to begin is with the claims of biologist M.T. Ghiselin (1974, 1981) and philosopher David Hull (e.g. 1978) about what biological species really are. To be members of the same species, individual animals must belong to historical lineages that have a common origin. They do not have to be similar to one another in any specified way. For example, there are no genes that every dog has in common with every other dog. Every dog gene has alleles. Similarly, there are no properties that every dog has in common with every other dog. Nor is it mere overlap in properties or resemblance to some paradigm that makes a group of dogs be conspecifics. Highly similar species but that have different historical origins do not form one species but several. Species, according to Ghiselin and Hull, are not similarity classes but big,

\* © Ruth Garrett Millikan 2005.

scattered, historical *individuals* enduring through time. They are entities somewhat like the Kennedy family, which is held together, of course, not by 'family resemblance' in Wittgenstein's sense but by blood-relatedness.

On the other hand, in the case of species, blood-relatedness is bound to be accompanied by considerable overlap in properties. If the species reproduces asexually, the reason is that progeny are clones. If the species reproduces sexually, then each of the genes in the gene pool has to fit in with a random selection of other genes from the pool so as to help produce a viable individual frequently enough to get itself reproduced often enough not to be eliminated from the gene pool. No single gene that changes the animal in very extreme ways can survive. This results in what is called 'homeostasis' in the gene pool. Thus, the various individuals within a species *mostly* resemble one another in a great variety of ways, but do not *all* resemble one another in any particular ways. But what pulls them together as a group is not just that they have common or overlapping properties, but that they tend to have common and overlapping properties *for a good reason*. One individual is like the next *for a good reason*. There is a good explanation of why one is likely to be like the next. Various kinds of inductions drawn over the members of a species are likely to be sound owing to certain kinds of causal connections among these members.

On this analysis of what pulls the members of a species together, species are not classes. Classes are defined by the members having certain common properties. Fuzzy classes may be defined by the members having overlapping properties or by their having many properties in common with a paradigm or paradigms. But the members of a class do not need to be like one another for any reason. They may be like one another quite by accident. Categories are classes or fuzzy classes. Species names are not names of categories.

Now, I need, first, to explain why this point is important and I need, second, to generalize it.

The point is important because it explains why it is possible to study a species as such, to gather stable information about it. If there is a reason why one dog is likely to be like the next in a good number of respects, then there is a reason why studying one dog is likely to yield

a considerable amount of probable knowledge about the next dog. In fact, of course, dogs are something it is possible to learn a great deal about. Consider how much time may be spent on learning about dogs by a student at veterinary college. True, all this knowledge is merely probable knowledge. Whatever one learns of the properties of dogs, it won't be analytic or necessary that every individual dog has each of those properties. But mere classes are not things one can learn anything at all about by induction. If there is no reason, given one member of a class, why the next member is likely to be like it, then if any inductions over the class turn out true conclusions, it can only be by accident. For example, it seems likely that there is no reason why one red triangular object should tend to be like the next in any respect other than redness and triangularity, so it is not likely that discovering, say, that one red triangular object is sweet will be of any use in predicting the taste of the next.

The way in which dogs are cemented into a unit is important, then, because it is only when individuals are cemented into a unit in some analogous way, such that there is a reason why one individual should be like the next, that we can obtain knowledge about this unit—unless, of course, by examining every member separately. It is obvious, then, why this sort of unit is the sort that tends to acquire a name. Names for mere classes are in most contexts quite useless. Names for units of this kind are names for the seeds on which all empirical knowledge is built, for all empirical knowledge is inductive.

The point about dogs is generalized by noticing other kinds of relations that tend to cement a unit together such that there is a reason why one individual within the unit or one part of the unit is liable to be like another. I have a name for units of this general kind, taken from Aristotle. I call them 'substances', a word which nonphilosophers may need to read as a new technical term but which philosophers may recognize as fairly traditional. Aristotle spoke of 'secondary substances'—the unit *dog* would be an example—and of 'primary substances', which were individuals. (Recall that what Ghiselin and Hull said about dogs was that they were big, scattered, historical *individuals* enduring through time.) There is a chapter in *OCCI* detailing many kinds of substances. Here I will discuss only a few, but enough, perhaps, to give you the flavor.

Substances fall roughly into at least three basic sorts, which I call 'historical kinds', 'eternal kinds', and 'individuals'. Historical kinds are like dogs. They are collections of individuals scattered over a definite spatiotemporal area that are causally related to one another in such a manner that each is likely to be like the next in a variety of respects. The two most obvious sorts of things that cause members of a historical kind to be like one another are these. First, something akin to reproduction or copying has been going on, all the various individuals having been produced from one another or from the same models. Second, the various members have been produced by, in, or in response to the very same ongoing historical environment; for example, in response to the presence of members of *other* ongoing historical kinds. A third and ubiquitous causal factor often supporting the first is that some 'function' is served by members of the kind, where 'function' is understood roughly in the biological sense as an effect raising the probability that its cause will be reproduced. It is typical for several of these kinds of causes to be combined. Artifacts are often good examples of this.

Consider chairs. Chairs have been designed to fit the physical dimensions and practical and aesthetic preferences of humans, who are much alike in relevant respects for the same reasons dogs are. Moreover, the design of a chair is pretty invariably influenced by the design of previous chairs, typically because these previous chairs have functioned well and were aesthetically pleasing within a cultural setting, relevant aspects of the cultural setting being reproduced elements as well. For these reasons, chairs form a rough historical kind. There are reasons that have nothing to do with any arbitrary points of definition why one knows roughly what to expect when someone offers to bring a chair.

Renditions of 'The Irish Washer Woman' or of *The Rite of Spring* form a historical kind. They are copied from one another or from scores that are transcribed from earlier renditions or copied from earlier scores. McDonald's restaurants form a historical kind. There are causes of their being so much alike. Professors, doctors, and businessmen form historical kinds, especially well-integrated ones when these groups are studied as limited to particular historical cultural contexts. Members of these groups are likely to act similarly in certain ways and to have attitudes in common as a result of similar training handed down from



person to person (reproduction or copying), as a result of custom (more copying), as a result of natural human dispositions (compare dog dispositions) or social pressures to conform to role models (copying again), and/or as a result of legal practices handed down from univocal sources. There is a reason why it may be productive to investigate, say, 'the attitude of American doctors towards acupuncture'. These attitudes are contagious. They spread.

The members of eternal kinds are like one another for a different kind of reason. They are alike because they possess a common inner nature of some sort, such as an inner molecular structure, from which the more superficial or easily observable properties of the kind's instances flow. The inner structure results in a certain selection of surface properties, or results in given selections of properties under given conditions. Popular examples of this sort of kind are the various chemical elements and compounds, along with various forms of these such as ice, liquid water, and steam. Portions of water have an inner structure in common that produces the same surface properties given the same temperature conditions. Strictly speaking, I suppose, gold, nitrous oxide, ice, and so forth are not kinds but stuffs, but samples of them are members of eternal kinds. Also, water molecules, electrons, protons, and so forth are examples of eternal kinds. Stars, planets, comets, asteroids, and geodes are eternal kinds, not because their properties flow always from exactly the same inner nature, but because they were formed by the same natural forces in the same sort of circumstances out of materials similar in relevant ways.

Eternal kinds can be said to have 'essences' in a very traditional sense, essences that are not nominal but real, often discovered only through empirical investigation. The reason that the members of these kinds have many properties in common is that they have a few fundamental properties and/or causes in common that account, given laws of nature, for all the others. Eternal kinds do form classes, all of whose members are alike in a variety of respects. But they are also much more than mere classes, because they are alike in these respects not by accident but in accordance with a causal explanation.

The last kind of (Aristotelian) substances are individuals. Individuals have been taken in modern times to have a very different sort of unity

than the unity of a kind, but there is a way in which the cement that holds a single individual together as it endures through time is quite a lot like the cement that holds a historical kind together. Ghiselin and Hull claimed that species are actually individuals, because they are held together not by a traditional essence but through historical causal connections. The other side of this coin is that individuals are rather like species. A species is a 'homeostatic system . . . amazingly well-buffered to resist change and maintain stability in the face of disturbing influences' (Eldredge and Gould 1972: 114; quoted in Hull 1978: 199). Similarly, an individual animal is a 'homeostatic system . . . amazingly well-buffered to resist change and maintain stability in the face of disturbing influences'. If a person is tall, brown-haired, knowledgeable about electronics, and a good piano player today, it is likely, though not certain, that she will have each of these traits also tomorrow. The various members of a species are like one another in part because they are, as it were, copied from one another. An individual physical object tends to have the same physical properties the next day as it had the day before because of natural conservation laws which tend not to copy, of course, but to preserve its properties from day to day. The effect, however, is much the same. The inference that an individual animal or inanimate object will probably have these and those properties tomorrow because it has them today is likely to yield a true conclusion for the same general sort of reason that an inference that other members of a species probably have these and those properties because this member has them is likely to yield a true conclusion. Individual objects are things that inductive knowledge can be collected about over time for the same sort of reason that historical kinds and, more broadly, eternal kinds are things that knowledge can be collected about over time.

I have explained why historical kinds, eternal kinds, and individuals, three basic kinds of (Aristotelian) substances, are similar with respect to the question why it is possible to gain inductive knowledge about one part of the cemented-together unity they compose from other parts. The reason this is possible for each is that it is not merely a class, either focused or fuzzy. Because substances are not classes, not units cemented together merely by some set of common or overlapping

properties, to have a concept of a certain substance is not to have a certain set of properties in mind, whether derived from paradigm cases or from exemplars. My next job then is to explain what it is to have a concept of a certain substance if not to have in mind a set of central properties. I am going to do this by explaining, first, what it is to have a concept of an Aristotelian primary substance—of an individual. Then I will generalize to other kinds of substances.

The idea that there are such things as ‘concepts’ of individuals is foreign to many psychologists and to many philosophers too. This is for interesting historical reasons that need not detain us here. If this use of the term ‘concept’ bothers you, then interpret me as just talking about thoughts of individuals or ideas of individuals. What is involved in being able to think of an individual?

One traditional twentieth-century answer to this question is that to think of an individual is to capture that individual with a description that uniquely identifies it. Another twentieth-century answer is that to think of an individual requires that you know how to identify it one way or another, perhaps by description and perhaps just by being able to recognize it, to differentiate it from other individuals, in perception. These views are close enough to the answer I would give myself that they will serve my purposes here. Something that they have in common and that I am sure is correct is the assumption that there is more than one way to think of the same individual; indeed, that there are innumerable ways to think of the same individual. An indefinite number of individuating descriptions apply to every individual. Similarly, there are, in general, numerous ways that the same individual might be recognized by sight, by characteristic sounds that they make, by smell (dogs are good at this), and so forth. Contrast the ways Helen Keller recognized her friends with the ways they recognized one another. Twentieth-century tradition had it, then, and I believe correctly, that there is no single or definite set of properties that one must either think of or be able to discriminate in order to have a concept, a thought, of an individual. Nor is there some central set of properties, some or most of which one must think of or be able to recognize in order to think of a particular individual. Similarly, I will soon claim, there is no central set of properties, all or some of which one must be able to think of,

recognize, or discriminate in order to think of the (Aristotelian) substance *dog*, hence in order to learn about dogs, to understand things said about dogs, and so forth. But I will come back to that part a bit later.

First, we have to deal with fallibility. The ways we have of recognizing individuals are always fallible in principle. Even supposedly individuating descriptions always presuppose that there is indeed one and only one thing fitting the description, something not guaranteed, for example, merely by the description containing superlatives. It might always be, for example, that no one is tallest or first in line, or first on the moon. More important, if you are actually to use an individuating description for purposes of recognizing an individual, you will have to recognize exemplifications of the properties mentioned in the description. But one's capacities to recognize objective properties are always fallible, for they depend on external intervening or mediating conditions such as lighting conditions, atmospheric conditions, sound absorption and reflectance properties of surrounding objects, obscuring conditions such as intervening objects, masking sounds and odors, and so forth. Nor is there an independent way of ascertaining what these mediating conditions happen to be in a particular case. There are always possible conditions under which you would misidentify or fail to recognize even your own mother or spouse.

Having the ability to recognize an individual, then, cannot be the same thing as being infallible in recognizing it. I have the ability to walk. It is one of my very best abilities. It does not follow that it cannot happen that I trip and fall when trying to walk. These reflections suggest that what we need here is an analysis of what it is to have an ability to do something, such as walking or recognizing your mother, that does not equate an ability with any simple sort of disposition. That analysis has been given in *OCCI*. But that all abilities are fallible is common sense, and I propose just to assume it here.

Tradition plus common sense suggests, then, not only that different people can have different kinds of concepts of the same individual by using quite different methods of recognition, but also that the methods any one person uses to recognize, hence to be able to think of, an individual will be fallible. Nor do these methods constitute a definition

of the individual. Your mother is not defined by the way you recognize her; say, by the look of her face and the sound of her voice. She doesn't have a definition, a set of properties, that make her be who she is. She is not a class that happens to contain only one member.

Similarly, the species *dog* is a unity that different people can have quite different kinds of concepts of, by using quite different methods of recognition. Whatever method a person uses for recognizing dogs, this method may always be fallible. Nor does the method that a person uses for recognizing dogs constitute a definition of what dogs are, even for that individual. The species *dog* is not just a class that happens to contain so-and-so many members.

What makes substances interesting is that there is often a great deal that can be found out and known about them. Often they have a great many properties. And it is typically the case that numerous of these properties and numerous sets of these properties will each be diagnostic of the substance. That is, each of these properties or property sets will be found only or typically when the substance itself is encountered. At least this will often be so within the spatial and temporal area inhabited by the person needing to recognize a substance. Mistakes that people might have made had they lived in different places and times are not relevant to their actual abilities to recognize substances. This is why it is possible for different people to have concepts of the very same substance by very different means. Children and chemists have different ways of recognizing sugar. You and Helen Keller have different ways of recognizing nearly every secondary substance, nearly every ordinary stuff, and nearly every ordinary eternal and historical kind.

Further, none of the ways that a person knows to diagnose the presence of a substance needs to be infallible. No particular set of properties used to diagnose a substance is ever definitional of it, although in the case, especially, of eternal kinds, empirical investigation may reveal (with probability) that, in fact, some sets are always correctly diagnostic. It is always logically possible that there is some other substance that has parts of its cemented-together unity that share the very same properties as the properties one is using, with practical success, for diagnosis of a certain substance. I can put this for

philosophers by saying that the possibility of ‘twinearth water’, certainly of ‘twinearth dogs’, and, indeed, of ‘twinearth Mama’, indistinguishable from your mother, is never ruled out by logic alone. It takes more than a set of properties in your mind to determine a substance. It takes a certain sort of causal glue in the world, holding that substance together. But given that glue in the world, conceptual access to that glued-together unity may be had by reference to any of many of its different parts or properties.<sup>1</sup>

In talking about what is involved in having a concept of a substance, I have quietly been making an assumption that I must now bring into the foreground as a claim. I have spoken of ways of recognizing a substance, and I have said that your ability to recognize a certain substance can depend on your inhabiting a certain space-time locale, one where certain diagnostic properties do mostly signify encounters with that substance rather than with others. The assumption I am making is that thinking of a substance involves the ability to recognize it, as it were, *in the flesh*, not merely the ability passively to contemplate its properties. We have thoughts of substances in order to be able to collect information about substances, which information we pick up on some occasions and then apply on others. To pick up information about a substance you must be in a position to interact with the substance, or with other things that interact with the substance, other things that are influenced by the substance or that influence it. Natural information is transmitted in the causal order, and you have to be in the causal order with whatever the information is information about in order to receive it.<sup>2</sup>

<sup>1</sup> Philosophers may detect a missing link in this analysis. The link is needed to connect the ability one has to recognize a particular substance to prior encounters with that particular substance rather than with similar substances on twinearth or wherever. That link is supplied in the description of abilities given in *OCCI*. What an ability is an ability to do is determined not merely by current dispositions but by the histories of the mechanisms responsible for those dispositions.

<sup>2</sup> I am using the notion *natural information* in a way somewhat like the way Dretske uses it in *Knowledge and the Flow of Information* (1980), yet not quite that way. For our purposes here the difference probably does not matter, but a careful description of the kind of information I have in mind is in Millikan (2004: chs. 3–5) where I call it ‘local information’.

Now, if you think about that claim for a moment, you will see that it is a fairly radical one. Surely you can have a concept of the last dinosaur species on earth to become completely extinct and of the first baby to be born next year, and of any other substance for which, although you have never encountered it, you do know an identifying description. And you have these concepts without having the slightest idea how to identify these things in the flesh. Surely you can have a concept of molybdenum—you can think about it and ask questions about it—without being able to identify it in the laboratory. Surely you can have a concept of Socrates without being able to identify him in the flesh, even if you were to be transported back to ancient Athens. Let me tackle the descriptions first, then come back to molybdenum and Socrates, for they will prove far more interesting.

The descriptions are handled this way. That your circumstances are such that you never get a chance to use an ability that you have does not take that ability away from you. You won't lose the ability to swim just because they chain you to a post in the middle of the Sahara desert for the rest of your life. If you understand the terms in any description and know how to apply them—that is, you know how to recognize the other objects and properties and relations mentioned in the description—and if you are right that the description is identifying, then you know a way to identify the substance that the description describes. You would do so by encountering something that you can recognize directly as fitting that description, or by coming across something else that you recognize as carrying information telling what fits that description. There are many cases in which you just aren't at all likely to come across any such information, but that is irrelevant to whether you have a capacity to recognize the substance. I am assuming here a fairly usual reading of the notion of natural information, according to which information about the past and about the future are entirely routine kinds of information (but see n. 2). And I am about to claim that language is a standard medium through which natural information is transmitted, hence a standard medium through which substances are recognized exactly as they are recognized 'in the flesh' through other media such as light and sound.

Now consider molybdenum and Socrates. It seems an obvious fact that many of our concepts of substances have been acquired without

encountering those substances directly but only by hearing about them. Moreover, as Kripke (1972), Putnam (1975), and Burge (1979) have observed, we often have no unique descriptions of these substances in mind either. How then can we be said to know how to recognize them? The answer, I claim, is that speech is just as direct a medium for the perception of objects and events and their properties as is the light reflected off objects, the smells emanating from objects, the sounds emanating from events in the environment, or the mechanical stimulations caused by objects in direct contact with one's body. This is a thesis that requires defense, and I have defended it at length both in *OCCI* (ch. 6) and in Millikan (2004: ch. 9). Here I can only throw out the rough idea, hoping that if it strikes you as dubious you will look to these longer versions and defenses before final judgment.

The claim is that hearing and immediately believing a sentence about a fact or occurrence is in relevant respects just like, for example, seeing that something is the case or seeing that something has occurred and immediately believing it. There is experimental evidence that what one is told goes directly into belief unless cognitive work is done to prevent this, just as what one perceives in other ways, through other media, does. Loading the cognitive systems with other tasks, such as having simultaneously to count backwards by threes, has the effect of facilitating belief-fixation regarding whatever one hears or reads (Gilbert 1993). Recognizing a linguistic reference to a substance is as much a way of recognizing the substance 'in the flesh' as any other way of recognizing it. It is identifying it and recognizing natural information concerning it through one more medium of manifestation. Think of this medium, the speech of another person, as like an instrument that aids perception. The lens of one's eye is, of course, an instrument that aids perception. If one wears corrective lenses, they are another such instrument. The speech of another person is analogous to somewhat more complicated instruments of this kind. Like a camera, a radio, a CAT scan, or a microscope, another person who talks to you picks up information-bearing patterns from his environment, focuses them, translates them into a new medium, and beams them at you. Think of living in a language community as like being inundated in one more sea of ambient energy. Like the surrounding light, surrounding people



transmit the structure of the environment to you in ways that, barring certain interferences, you have become tuned to interpret. Becoming tuned to interpret the information-bearing patterns that are common in a certain language community is coming to understand the language of that community. Similarly, a radiologist must learn to interpret the information contained on X-ray images and the auto mechanic must learn to interpret the information contained in the sounds emanating from ailing automobile engines.

The notion that understanding and believing what is said to you is just one more level of natural-sign reading on the same level as ordinary perception is to many people quite unintuitive. One reason is that what is given to you in ordinary perception is always given as in some quite definite current relation to you. It is given as happening at the time you perceive it, as happening relatively nearby, and often as bearing quite an exact spatial relation to you. This kind of information is needed to guide action, for how one can currently act on a thing always depends on its current relation to oneself. Ordinary perception is for immediate action, whereas what one learns through language is not typically used that way. Usually you are not told what exact spatial and temporal relations the objects and events being presented to you through language have to you here and now. But there are intermediate cases; for example, video recordings. It is clear enough that you *perceive* things happening when you watch a video, but, as in the case of language understanding, you do not perceive the spatial and temporal location to yourself of what occurs on video.

A second reason that the comparison between ordinary perception and language comprehension is unintuitive is that ordinary perception is so much more reliable than what one hears said, at least under common circumstances. It is not easy to fool ordinary perception. To create strong perceptual illusions requires a good deal of knowledge about the perceptual mechanisms and often quite special equipment, of the kind, for example, that optometrists have in their examination rooms. This is a difference of degree, however, a mere difference in frequency, not a difference in kind. Recalling that film dubbing is currently the rule rather than the exception, what differences are there, for example, among (1) believing what you apparently see when

you look through the peephole into an Ames room, (2) believing what you see when a film has been dubbed, and (3) believing what you hear someone say when it's false? In the modern world, if you want to believe only what's true, you often have to apply heavy filters to other methods of perception as well as to perception through language.

The upshot of these reflections is that we can understand how it is possible to recognize a substance through the information that language bears; indeed, how it is possible to come to be able to recognize a substance pretty much merely by learning a word for it. This is how we manage to have concepts of Socrates and, for most of us, how we manage to have concepts of molybdenum. To have a word for a substance is to have an essential part of an ability to recognize manifestations of it that are generated in a particular language community. That, I have argued, is why it is possible for small children to learn, as Chomsky puts it (1995: 15), 'a word an hour' between 18 months and 6 years of age.

## REFERENCES

- Burge, T. (1979), 'Individualism and the Mental', in P. French, T. Uehling, and H. Wettstein (eds.), *Studies in Metaphysics: Midwest Studies in Philosophy, iv* (Minneapolis, Minn.: University of Minnesota), 73–121.
- Chomsky, N. (1995), 'Language and Nature', *Mind*, 104: 1–61.
- Dretske, F. (1980), *Knowledge and the Flow of Information* (Cambridge, Mass.: MIT Press).
- Eldredge, N., and Gould, S. J. (1972), 'Punctuated Equilibria: An Alternative to Phyletic Gradualism', in T. J. M. Schopf (ed.), *Models in Paleobiology* (San Francisco, Calif.: Freeman, Cooper), 82–155; quoted in Hull (1978).
- Ghiselin, M. (1974), 'A Radical Solution to the Species Problem', *Systematic Zoology*, 23: 536–44.
- (1981), 'Categories, Life and Thinking', *Behavioral and Brain Sciences*, 4: 269–83.
- Gilbert, D. (1993), 'The Assent of man: Mental Representation and the Control of Belief', in D. M. Wegner and J. W. Pennebaker (eds.), *Handbook of Mental Control* (Englewood Cliffs, NJ: Prentice-Hall).
- Hull, D. L. (1978), 'A Matter of Individuality', *Philosophy of Science*, 45: 335–60; repr. in Sober (1994: 193–216).

- Kripke, S. (1972), *Naming and Necessity* (Cambridge, Mass.: Harvard University Press).
- Millikan, R. (2000), *On Clear and Confused Ideas: An Essay on Substance Concepts* (Cambridge [OCCI]: Cambridge University Press).
- (2004), *Varieties of Meaning* (Cambridge, Mass.: MIT Press).
- Putnam, H. (1975), 'The Meaning of "meaning" ', in Keith Gunderson (ed.), *Language, Mind and Knowledge; Minnesota Studies in the Philosophy of Science*, vii. (Minneapolis, Minn.: University of Minnesota Press).
- Sober, E. (1994) (ed.), *Conceptual Issues in Evolutionary Biology* (Cambridge Mass.: MIT Press).

## 7 Cutting Philosophy of Language Down to Size\*

---

When asked to contribute to the 'Philosophy at the New Millennium' lecture series, my first thought was to talk about philosophy of biology, a new and increasingly influential field in philosophy, surely destined to have great impact in the coming years. But when a preliminary schedule for the series was circulated, I noticed that no one was speaking on language. Given the hegemony of philosophy of language at mid-century, after 'the linguistic turn', this seemed to require comment. How did philosophy of language achieve such status at mid-century, and why is it losing it now? Has the Anglo-American tradition really begun to put the philosophy of language in better perspective? I hope so. Indeed, I will end with suggestions for how to keep it more securely in its proper place.

William Alston is one of the most respected and well-balanced philosophers of language from the period of its heyday, and we are blessed with a characteristically lucid article of his on philosophy of language written for the 1967 edition of the *Encyclopedia of Philosophy* in which he explains why this subject is so central (vol. iv, pp. 386–90). 'Whenever philosophers turn their attention to any subject matter . . . one thing they try to do is to clarify the concept of that subject matter.' Accordingly, he says, a view that many philosophers came to hold was that the primary job of the philosopher is 'conceptual analysis'. Alston then claims that a 'shift in the centre of gravity of philosophy' took place with a shift in understanding of

\* First published in A. O'Hear (ed.), *Philosophy at the New Millennium*, Royal Institute of Philosophy Supplementary Series (Cambridge: Cambridge University Press, 2001), 125–40.

what 'conceptual analysis' is. Rather than being understood as an attempt directly to analyse the nature of its subject matters, for example of *causality* or *knowledge*, or as an attempt directly to explore these concepts as psychological phenomena, the conceptual analyst was now understood as 'trying to make explicit what it is one is saying when he says . . . '—this or that pertaining to *causality* or *knowledge*. Although Alston doesn't remark on it, that, of course, was how this tradition of philosophy managed to continue to understand itself as an a priori discipline after the demise of rationalism. Otherwise it was thought unclear why philosophers wouldn't need to turn to data or to experimental work in order to carry out their tasks.

Alston proceeds:

In so far as conceptual analysis takes the centre of the stage the pursuit of the philosophy of language receives a powerful impetus . . . the analytic philosopher will spend a great deal of his time in trying to decide what a given expression means, whether two expressions mean the same thing, how a given expression is used, and so on.

Thus,

the need arises to develop an explicit theory about what it is for a linguistic expression to have a certain meaning, what it is for two expressions to have the same meaning, and what it is for a sentence to be used in order to do so-and-so . . . it follows that in so far as philosophy consists of conceptual analysis it always deals with language in one way or another. This might be taken to imply that philosophy of language is the whole of philosophy and one might choose to use the term 'philosophy of language' in such a way that this would follow.

Alston does not recommend this usage. But it is clear that the conclusion could easily be drawn that philosophy of language has something like the status of first philosophy, since in an important way all other philosophy depends on it. Moreover, it seems to follow, paradoxically, that by underpinning all of philosophy it also underpins itself. For what is the methodology of the philosophy of language itself to be? Presumably, it too will proceed by 'trying to make explicit what it is one is saying when he says . . . ' ; say, that a word has a certain referent or is synonymous with another word.

To appropriate the methodology of conceptual analysis as basic is, I believe, implicitly to make certain very strong assumptions about the nature of language, which assumptions cannot, in principle, be questioned by use of the method. Then, it is by begging the question that the philosophy of language promotes itself to become the foundation for the entire philosophical enterprise. These two strong assumptions concern referential or extensional meaning. (I will just say 'referential', terminologically counting properties and relations as 'referents'.) They are assumptions that externalist theories of thought content and direct reference accounts of word meaning have been trying to challenge for some years. But neither of these latter two programs can be carried through, I believe, without a fundamental shift in methodology, a shift that amounts in the end to abandonment of the position that philosophy is always, at root, an a priori discipline.

The first assumption I call the 'seed assumption'. Unlike the act of *knowing* an empirical fact, which act can be successfully completed only through cooperation of the external world, the act of *referentially meaning* something, though perhaps not of actually referring to something (recall Frege's empty descriptions) is completed within the mind itself. The intentional nature of the act of referring has its source or is given its shape by the mind. The mind, or its contents, alone determines the *criteria* in accordance with which a reference succeeds or fails to be made. The seeds of reference, if not the flower, are always entirely within the mind.

The second assumption I call the 'one-to-one assumption'. A univocal term in a public language is associated with one psychological state common to all competent users. For referential terms, typically the idea has been that the same seed of reference or criterion for successful reference must be grasped by all of its competent users. The paradigm here is the Fregean view that an unambiguous word or sentence corresponds to a single *Sinn* that is grasped by all who understand it.

The one-to-one assumption explains why conceptual analysis can be applied to referential terms in a public language. If there is something common to the psychological states of all who grasp the meaning of a public term, then there will be some single correct description that captures this common element, and any competent user should be able

to recognize that description. The one-to-one assumption taken with the seed assumption leads to the view that the various types of referring or denoting expressions—proper names, common names, names of properties and relations, definite descriptions, indexicals—are mirrored in distinct types of mental structures. Then there is a one-to-one relation not only between referential linguistic meanings and inner seeds of meaning, but also between types of referring expressions and types of meaning seeds. These two assumptions are independent. The one-to-one assumption can be held without the seed assumption; for example, by taking linguistic meanings to be stereotypes that don't determine extensions. And the seed assumption can be held without the one-to-one assumption. There might be such things as inner seeds of referential meaning and yet public communication might characteristically rest only on shared reference. This last position has tempted some, for example, in the case of proper names.

The program of conceptual analysis is most agreeable when coupled with some form of linguistic idealism. If our language determines our concepts and categories and these in turn determine the forms we use to mould and measure our world, then to explore the structure of our concepts and categories will be to explore the structure of our world. To project the same conceptual schemes as one another will be to live in the same world. Thus, for example, the idea that we are the authors of the criteria of identity used to individuate various kinds of objects—that terms such as 'apple' and 'rabbit' 'possess built-in modes, however arbitrary, of dividing their reference' (Quine 1960: 91)—clearly entails that the referents of our words and thoughts are determined to fit structures that we project. An examination of our peculiar cognitive ways will reveal the structures that our referents are necessarily determined to have. Assume that these structures are acquired by internalizing the rules of a public language, and the result is an extreme form of linguistic idealism. Our language constitutes our world. But then it must also constitute *itself* within that world since, of course, it is part of the world. Philosophy of language is called on to support itself as well as the rest of philosophy.

Less radically, concepts may be understood to be nodes in a net of response and inference dispositions, their references determined by

their positions in the net plus, perhaps, their dispositional causal relations to objects and properties in the world—their ‘causal roles’. (The classical notion that concepts correspond to sets of necessary and sufficient conditions for application and also the view that concepts are organized around prototypical instances are, of course, merely simple variations of this theme.) Then, to analyze the position of the concept in the net will be to trace out central requirements on the nature of its referent. Assume, again, that these inference dispositions are acquired by internalizing the rules of a public language, and it again follows that the meanings of referential terms can be traced by conceptual analyses. The principle is the same, of course, if the rules or habits of language are conceived as governing rich ‘forms of life’ and not merely perceptual judgments and inferences. Quine, Wittgenstein, Sellars, and so forth, all look very much alike when considered in these general terms, as do nearly all contemporary psychologists studying concepts. All make both the seed assumption and the one-to-one assumption.

Two important challenges to these assumptions are familiar to us all. First is the effort to evict ‘meaning’ from the ‘head’ along with the defence of ‘direct-reference’ (Quine 1953: 20–46) views, begun a quarter of a century ago by Putnam (1978*b*) and Kripke (1972). The denial that ‘meaning is in the head’ looks as though it should threaten the assumption that the seed of referential meaning is wholly in the mind. Second is the attack on the analytic–synthetic distinction and on the notion of linguistic meaning spearheaded half a century ago by Quine. These latter moves look as though they should threaten the assumption of a one-to-one relation between public linguistic meanings and psychological states. It looks as though no inseparable inference connections will be learned as one learns one’s language. Between them these two challenges have considerably diminished the influence of runaway conceptual analysis, but neither has succeeded in extracting its roots. I would like to understand the reasons for this failure.

Because the one-to-one assumption is so commonly made, the literature is often unclear about whether it is linguistic meanings that are being examined or merely referential thoughts. I will try to keep



these issues entirely separate, discussing, first, only the seed assumption. Later, I will discuss the one-to-one assumption.

Putnam (1975*b*; also in 1975*a*) is credited with launching the 'externalist' position that 'meaning is not in the head'. What Putnam claimed, more exactly, was that knowing the meaning of a natural-kind term was not merely a matter of being in a certain psychological state. Putnam's claim, as it bears on the seed assumption, was, for example, that if the psychological state associated, for any one of us, with the term 'water' were to be realized in a person living on twinearth, this psychological state would constitute a thought not of water but of twin-water. Putnam contrasted his position with the classical one that the extensions of thoughts of natural kinds are determined via prior thoughts of their defining properties. Instead, thoughts of natural kinds are, he said, 'indexical'. Facts about the environment of the thinker help to determine the thought's extension. Putnam's position has been considered a kind of 'externalism' concerning content for natural-kind terms, 'externalism' being touted as new and quite radical. Putnam himself compared his position to Kripke's position, also taken to be new and radical. He said that natural-kind terms, like Kripke's proper names, were rigid designators, not standing in each possible world for whatever extension certain properties have in that world, but always standing for the same kind, determined by its disposition relative to the thinker in *this* world. But I'd like to compare Putnam's move with one made much earlier in the century by Bertrand Russell concerning thoughts of individuals.

F. H. Bradley had worried that to make a judgment about a particular individual one would have to grasp its entire complex individual nature, so as to distinguish it in thought from all other possible individuals. But Russell said no. All you need is to distinguish the individual from other *actual* individuals with a true definite description. However, such a description would usually have to begin with a reference to the individuals' relation to some prior individual or individuals, say to Socrates or to Greece or to this earth, the threatened regress ending with a final reference to oneself. For example, I might end by making references to the causes of various of my current sense data—*this* one, *that* one—thus locating the object of my thought in relation to me in the historical (the

causal/spatiotemporal) order. Or I might think of James as whoever bears the relation 'same person' to whatever causes the characteristic such-and-such- (Jamesish-) looking and such-and- such- (Jamesish-) sounding sense data around here.<sup>1</sup>

The important point is that, on this sort of analysis, you have to be in the same world as the individual you are judging about, existentially<sup>2</sup> related to it in the right way. The possibility of having thoughts of individuals rests on the actual historical relations they bear to you, or, putting this broadly, on the actual-world context of your thinking. To my knowledge, no one after Russell has questioned this view. What individual my thought is about is determined by existential relations external to my mind. Another person in exactly the same psychological state might well be thinking of a different individual.

Does it follow that meaning has not been thought to be 'in the head' in the case of thoughts of *individuals* for nearly a century?

Surely it was not meaning, but merely extension, that Russell discharged from the head in this way. Whether we separate the meaning from the denotation (Russell's way) or the *Sinn* from the *Bedeutung* (Frege's way) or the intension from the extension (Carnap's way) or the meaning from the reference (Quine's way) makes little difference. The division has always been between what was fully determined within or before the mind, and what was partly determined by (in a broad sense) the external context of the mind. Meaning has been whatever is left in the head after you carve off those parts of the completed act of referring that lie outside. And it has always been meaning alone that determines what properties or relations confirm a certain extension as the one

<sup>1</sup> Of course, Russell took definite descriptions to express disguised judgments, so that this way of judging about an individual did not involve a thought whose referent was that individual. The judgment really concerned the contents of the world entire—that given the world entire, certain propositional functions were or were not always true. The disagreement between Frege and Russell over whether judging about individuals is done with an existence claim or with a truly referential thought but that might not have a referent is trivial in this context. When we consider thoughts rather than the functions of language forms, what is the difference between a presupposition of existence and a belief in existence? The difference seems merely a matter of what one is psychologically prepared at the moment to bring into question.

<sup>2</sup> By 'existential' I mean within the historical order. 'Existential relations' are spatial, temporal, or causal relations.

referred to. What's inside determines how things must be outside for the reference to be successful. No matter how much external-world involvement there is in reference, the seeds of referential meaning remain in the mind. That is the classical position.

So now Putnam tells us that there are mind-independent natural kinds out in the world waiting to be named by us just as there are individuals. And just as with individuals, you don't have to think of them by knowing their essences. You need not distinguish them in thought from all other logically possible kinds, but only from the other actual kinds in this world. Just as with individuals, you can point to them in thought with descriptions containing indexicals, your thoughts completed through existential relations to them. Water is whatever bears the relation 'same liquid' to the stuff in the lakes and streams *around here*, or whatever bears the relation 'same liquid' to the stuff causing and such (the waterish) surface properties *around here*. Nothing is new from the standpoint of the philosophy of language and mind. Meaning remains as much in the head as it was. Only the realist ontology is new, the (re)introduction of kinds that are not merely nominal. And even here Putnam did not argue that the criterion of sameness for 'same liquid' is nature-given rather than nominal. The conceptual analysts' move then is perfectly obvious. Treat Putnam's analysis as introducing definite descriptions in place of definitions or graspings of essences. Nor has there been a shortage of variations on this theme.<sup>3</sup> Certainly, Putnam was never clear on why this obvious move would not be possible, on why meaning needed to be pushed out of the head.

Similarly, Putnam's theory of the 'Division of Linguistic Labor' (1975*b*) reduces to the claim that some people think of some kinds by way of descriptions such as 'the substance the chemists have in mind when they say "molybdenum" '. What is threatened here is not the assumption that meanings are in the head, but rather the one-to-one assumption, the assumption that everyone who understands the linguistic meaning of a public referential term must think of its referent in the same way. Burge's claim (1979) that what a person means may

<sup>3</sup> For a very clear essay on this theme see Fumerton (1989).

depend on the habits of the surrounding language community can be trivialized in a similar manner. Kripke's suggestion on the reference of proper names becomes just (here I am caricaturing) that people not present at Aristotle's baptismal ceremony describe him to themselves as 'the person at the beginning of a chain of "Aristotle" tokens, certain more recent members of which chain are *this* token and *that* token'.

To clinch this deflationary account, we call attention to the method by which philosophers have arrived at their so-called externalist theories of thought content. They have used the method of conceptual analysis. They have brought forward examples and counterexamples to be judged for their conformity with our intuitive notions of what the referent of a person's thought would be, given this or that circumstance.

Some externalists have tried trading in descriptions for indexical 'characters', in Kaplan's sense. Rather than containing a description of the relation its referent must bear to the thinker, the referential thought is purely indexical. It inarticulately points to that which bears a certain existential relation to it. Compare a token of the word 'I'. It refers to whoever has produced it, but it does not describe this relation. It does not mean 'the person who produced this token'. But indexicals are semantic elements too. They have meaning only by a prior assignment of character. What then gives a mental indexical its character? What determines which particular existential relation its referent must bear to it? The obvious answer is given by Ned Block (1986). The rules determining the values of mental indexicals must be implicit in the cognitive dispositions of the thinker, in what the thinker would ultimately be disposed to identify as her referent given these or those, perhaps unexpected, circumstances. The seed of referential meaning remains in the head.

The boldest externalists are 'direct-reference' theorists, holding that some or all of our referential thoughts are 'Millian'. No semantic analysis can be given of a Millian thought beyond saying what referent it brutally names. But then a theory is needed telling what *constitutes* that one thought brutally has one referent while another thought brutally has another. Various theories of thought content have aspired to this, such as Dretske's informational theory, Fodor's asymmetrical

dependency theory, Papineau's or my teleological theories. Each tells its own story about the sort of existential relation that cements a mental representation to its referent. But how do we prove that one of these relations rather than another is the reference relation? If asymmetrical dependency were the reference relation, for example, then of course referential meaning would be fully external. Similarly, if being the last thing you stepped on before that thought occurred to you were the reference relation, again, referential meaning would be fully external. If you think the former more plausible than the latter (marginally!), how can you prove it?—unless, as Putnam put it (1978*b*), you can show that the former is more like the 'intended' relation of reference, more in accord with what we mean by 'reference'? Philosophy of language supports itself again! Putnam called it 'internal realism'.

Think of the matter this way. How could something be the criterion determining the referent of my own term or thought, if what I actually use is another criterion? Of course, I may have in reserve some more ultimate criteria that would prompt me to readjust my current criterion, for example, if I should find out more about the inner nature of the liquid I call 'water', or if I should find out how certain experts distinguish molybdenum. But these adjusted criteria would still be *my* criteria, set by my criteria for setting criteria. My meaning remains in my head. If the externalist proposes that something is 'reference' that we do not recognize by our own lights as reference, then he is just changing the subject. No way, it seems, can we defend any fully externalist view of referential meaning.

How do we exit this loop? This is what I think is needed. First, construct an externalist theory of reference that portrays it playing an indispensable role in the uses of both language and thought. Second, given this theory, make plausible that reasonable people should have come to the reasonable but mistaken conclusion that references/extensions are determined by criteria projected by individual minds. Third, make plausible that, despite their corresponding to different conceptions, the reference described by the theory is the same as the reference of our ordinary term 'reference'.

These can all be accomplished, I believe, though only the briefest sketch is possible here.

The inner-seed assumption, expressed in Putnam's terms, is that the referent of a thought must always be something 'intended' by the thinker. Put this a slightly different way: It must have been a purpose of the thought to grasp that referent. Stated this second way, I believe the assumption is correct. The shift is from a purpose of the thinker, a psychological purpose, to a purpose of the thought, a biological purpose. The purposes of my thought need not be purposes of *mine*. But it is certainly understandable that classical theories of reference should have confused these two kinds of purposes. This is why thoroughly reasonable people have been led to adopt the seed assumption.

Consider, for example, the purpose of my breathing. The purpose of my breathing is not a purpose of mine, not determined by an intention of mine. A frog's breathing has exactly the same purpose as mine, determined in the same way—certainly not by the frog's intentions. Similarly, neither the purposes of the frog's vision and hearing nor the purposes of mine are determined by anyone's intentions. Why then should the purpose of my thoughts, in particular of my referential thoughts, be determined by my intentions? More plausibly, the purposes of all of these are biological, not psychological, purposes.

Now, whatever gloss you put on the notion 'biological purpose', or 'biological function', one thing is clear. Biological purposes are not always achieved. In the biological world, failure in the particular case may even be the rule rather than the exception. The vast majority of individual animals die before reproducing. It would be very surprising if the biological purposes of human thought were invariably achieved. It is plausible that the various mechanisms designed to aid in forming empirical concepts are more successful sometimes than others. However, just as it may be perfectly clear what the goal was of the cat that pounced but missed the mouse, sometimes it may be perfectly clear what the goal of the concept-forming mechanisms is, even though they are not currently dead on target.

The central task of empirical concept-formation, I have argued, is learning to recognize what is objectively the same through enormous diversity of appearances, a wide variety of proximal stimulations (Millikan 2000). The cognitive systems succeed or fail in the

performance of quite specific tasks, such as the task of learning to recognize Mama reliably, or dogs, or orange juice, or doorknobs. These specific tasks get set as goals by the system when it is placed in a certain environment and supplied with a certain history of experience. Just as a dog that is tracking a rabbit is tracking some *particular* rabbit, though of course it may sometimes lose the scent, when the cognitive systems attempt reidentification of something, perhaps using certain marks or criteria as a guide, whether they are successful or not, usually there is something entirely definite and real it is their current purpose to be tracking and to be learning better ways to track. Thus, reference can remain stable while ways of identifying change, both during individual cognitive development and during the history of science.<sup>4</sup>

So it is possible coherently to question the inner seeds of reference assumption, the first assumption underlying the method of conceptual analysis. The second or one-to-one assumption, that the linguistic meanings of referential terms correspond to particular psychological states, is seldom discussed, even though numerous difficulties resulting from it are well known.

Starting with Russell, indeed even with Frege, it was recognized that if proper names abbreviated descriptions, still nothing seemed to prevent different speakers from using different descriptions to back the same name. Nor is the problem specifically with description theories. Insofar as thinking of an individual depends on a unique existential relation one has to that individual, each of us will think of the same individual in a different way. Suppose, for example, that for each of us—by description, or indexically, or causally—Aristotle is the man at the beginning of a causal chain issuing finally in this particular token and that particular token having the shape or sound ‘Aristotle’. Still, each of us will begin by confronting different tokens of ‘Aristotle’, or if the same token, then it will be related to us in different ways. If we focus instead on nondescriptive ways of recognizing, the scattering is even more obvious. I can recognize my daughters Aino Millikan and Natasha Millikan in hundreds of different ways. Now that you’ve heard of them, you can use their names too, but are

<sup>4</sup> For much fuller detail see Millikan (2000).

you able to recognize them in any of my ways? True, both of us now know them as my daughters. But many of their acquaintances, who surely mean the same as we do when using their names, have no idea who their mother is. That there must be some inner psychological state common to all who comprehendingly use a proper name seems entirely out of the question.

Turning to names for kinds or properties, the valid conclusion of Putnam's arguments for division of linguistic labour, and also Burge's arguments against 'individualism', is not that meaning is not in the head. Meaning-in-the-head can always be saved by putting more complicated descriptions in the head. What these arguments *do* show is that people who use a natural-kind term such as 'molybdenum' or 'arthritis' think of the same natural kind using different methods, some relying on experts, some not. And difference in method goes much deeper. Suppose in the case of the familiar stuff water, for example, that we all think of it, simply, as the watery liquid around here. Still, why suppose that the property *watery* is recognized in the same way by all speakers of English? Might it be that Helen Keller actually *understood* the word 'water' that Annie Sullivan spelled into her hand? Putnam says that the ways people recognized water prior to discovering that water is H<sub>2</sub>O did not determine the extension of their thoughts of water, but were subject to replacement by more accurate measures. Do today's children fail to understand the contemporary term 'water' before they have learned chemistry? What in the conception of water has to remain the same to prevent the word 'water' from changing its meaning?

If there were, indeed, something in common among the psychological states—the methods of recognition, inference patterns, or whatever—of all who understood a univocal referential term, how would these states be uniformly acquired in the course of language learning? Quine's arguments in *Word and Object* quite effectively showed, I believe, that for the most part there could not be such a thing as the inner rules of a language, the correct criteria for application or methods of recognition, the correct entailments, and so forth, that constitute public linguistic meanings. For the most part, everyone has his own private 'sentence associations', his own private application and inference procedures. Each follows his own causal pathways. There is



no way to enforce any particular ones of these to be constitutive of the language being learned. Again, Helen Keller both wrote and spoke English.

Wittgenstein suggested that the criterion for meaning the same was agreement in judgments. *Why* do we think that Helen Keller spoke English? Because of agreement in judgments. Those around her acknowledged, for the most part, the same mundane truths that she did. She merely discovered these truths in a different way. But agreement in judgments gives evidence only of agreement in reference. It tells nothing about the private methods by which those judgments were reached. Yet agreement in judgments, for the most part, is the only evidence we have when trying to coordinate our language use with others. This is not incompatible with there being *some* judgments you can learn to agree on only by following a specific pattern of recognition. If you are going to manage to agree in judgments about bachelors you must begin by having some way of recognizing incoming information that concerns men as such, and a way of recognizing incoming information that concerns marriage. That this is the only way to manage agreement in the case of judgments about bachelors is not a matter of the rules of English. It is a matter of the constitution of bachelorhood. There just aren't any other reliable signs of this state to go on (Millikan 2000: sect. 3, 4).

Yet there has been a dogged persistence in the view that a univocal linguistic meaning corresponds to a single psychological state. Thus, Quine, rather than conclude that the public meaning of referential terms might be, simply, their referring, proclaimed the 'indeterminacy of translation'. He *trashed* linguistic meaning! Putnam's reaction was equally bizarre. After concluding that linguistic meaning is in part a function of reference, he anxiously offered a substitute internal sort of meaning that competent speakers of a language were to be required to share:<sup>5</sup>

Speakers are required to know something about (stereotypical) tigers in order to count as having acquired the word 'tiger' . . . After all we do not permit

<sup>5</sup> True, he uneasily speaks here of 'acquiring words' rather than of learning their meanings, but the shift in terminology is not material.

people to drive on the highways without first passing some tests to determine that they have a *minimum* level of competence . . . English speakers are required by their linguistic community to be able to tell tigers from leopards . . . this could easily have been different (1975*b*: 168)

A more flimsy fantasy has never been concocted, I submit, by a competent philosopher.

There are, of course, classical arguments—four of them—that referential terms must have, besides public reference, also a public sense. One concerns assertions of identity, the second, assertions of existence, the third, interchangeability in intensional contexts, and the fourth, psychological explanation.

First, if only reference is shared by everyone for a referential term, how can an assertion of identity be informative? Peter Strawson provided the answer in 1975. He explained that what an identity sentence,  $A = B$ , does to your head is, as Michael Lockwood rephrased it (1971), to merge the contents of two information folders, one of which formerly had been accessed with the word 'A', the other with the word 'B', so that all the information can now be accessed with either word. Why wasn't this the end of the matter?<sup>6</sup>

Second, if a referential term, 'A', has no univocal public sense, what meaning will the assertion 'A exists' have, or the assertion 'A doesn't exist'? In the same spirit as Strawson, we answer that what 'A doesn't exist' does to your head is remove your mental folder of virtual information about A to an inactive file, such that it does not engage in the normal way with perception or action. 'A' will now be treated as only a pretend name. 'A exists' reverses this effect.<sup>7</sup>

Third, if 'A' and 'B' name the same, why is 'x believes that A is A' not always equivalent to 'x believes that A is B'? Because, very occasionally, in a context that makes this clear, using a certain word to describe the content of a person's belief carries the implication that the person himself would accept this word as helping to express this belief. But a person might believe that A is A without knowing that A is called 'B' (Millikan 2004: ch. 7).

<sup>6</sup> See Millikan (1984: ch. 12).

<sup>7</sup> Ibid.

Fourth, we describe the contents of people's beliefs and desires using embedded sentences, and having captured their contents in this way, we successfully predict their behaviors. How is this possible if the embedded sentences don't correspond to determinate psychological states? The answer is that this whole image of how behaviors are predicted is quite wrong. Comparative psychologists learned long ago that only the distal and not the proximal behaviors of an animal are predictable. You can predict that the cat will move the correct lever in the right direction to let itself out of the experimental box, but you can't predict whether it will use the left paw or the right, or perhaps lean with its tail. Having learned which way the lever needs to be pushed, it will do this in different ways at different times. Similarly, you can predict that Sam, who sincerely said he would arrive at seven, probably will, in one way or another—by bus, by subway, by car, on foot or bicycle—arrive around seven. And if he says he will take the 6 o'clock bus, you can predict that he will probably manage to identify the 6 o'clock bus. There is no need for concern with his methods of recognition. True, occasionally it is helpful to know something about the ways a person can or can't identify a certain referent. If you wish to know why Sam didn't even attempt to speak to his favourite author who was at the same party he was, it may help to know that Sam doesn't know what she looks like. But if he does know what she looks like, it is not likely to be useful to know whether it's the line of the nose or the chin that cues him. Similarly, when you order basil with your groceries, you are not concerned whether the grocer recognizes basil by the smell, or the shape of the leaf, or the label on the box, so long as it's basil that he brings.

Notice that the attempt to understand the meanings of words and the reference of thought in this way, without the seed assumption and without the one-to-one assumption, besides requiring development of certain broad aspects of a realist ontology, also implicates broad issues in theoretical cognitive psychology, child development, and, indeed, the history of science. Understanding what reference is is not conceived as a wholly a priori project. The aim is to help create the framework for an empirical theory of human cognition. Indeed, if the nature of reference is at all as I have suggested, then to investigate any of our

empirical concepts is to investigate the nature of the world, not merely of what's in our heads. This is consonant, I believe, with the actual role philosophy has always played in relation to the sciences. The philosophies of mind and of language, for example, have provided the foundation for much that has happened in modern empirical psychology and linguistics. Large parts of philosophy have traditionally operated as theoretical empirical sciences. There is also a long tradition, including, for example, Russell, that has explicitly endorsed this kind of role for philosophy.

Dropping the seed assumption and the one-to-one assumption puts a very different face on both the appropriate methods in philosophy of language and its status relative to other parts of philosophy and science. I hope this is the face it will wear in the coming century.<sup>8</sup>

## REFERENCES

- Block, N. (1986), 'Advertisement for a Semantics for Psychology', *Midwest Studies in Philosophy*, 10: 615–78.
- Burge, T. (1979), 'Individualism and the Mental', in P. French, T. Mehling, and H. Wettstein (eds.), *Studies in Metaphysics: Midwest Studies in Philosophy*, iv (Minneapolis, Minn.: University of Minnesota), 73–121.
- Fumerton, R. (1989), 'Russelling Causal Theories', C. W. Savage and C. A. Anderson (eds.), in *Rereading Russell: Essays in Bertrand Russell's Metaphysics and Epistemology*, *Minnesota Studies in the Philosophy of Science* xii (Minneapolis, Minn.: University of Minnesota Press), 108–18.
- Hempel, C. G. (1950/65), 'Empiricist Criteria of Cognitive Significance: Problems and Changes', repr. in *Aspects of Scientific Explanation and Other Essays in the Philosophy of Science* (New York: Free Press Macmillan), 101–22.
- Kripke, S. (1972), *Naming and Necessity* (Cambridge, Mass.: Harvard University Press).
- Lockwood, M. (1971), 'Identity and Reference', in M. K. Munitz (ed.), *Identity and Individuation* (New York: New York University Press), 199–211.

<sup>8</sup> It will surely have been noticed that this essay has presupposed a heavy-duty realist ontology throughout. The ontology is discussed at length in Millikan (1984: chs. 16–17; 2000: chs. 2–3).

- Millikan, R. G. (1984), *Language, Thought, and Other Biological Categories* (Cambridge, Mass.: MIT Press).
- (2000), *On Clear and Confused Ideas* (Cambridge: Cambridge University Press).
- *Varieties of Meaning* (Cambridge, Mass.: MIT Press).
- Putnam, H. (1975a), *Mind, Language and Reality: Philosophical Papers, i*, (Cambridge: Cambridge University Press).
- (1975b), 'The Meaning of "Meaning" ', in Keith Gunderson (ed.), *Minnesota Studies in the Philosophy of Science vii*, (Minneapolis, Minn.: University of Minnesota Press). (Also in Putnam, *Mind, Language and Reality* (1975a)).
- (1978a), *Meaning and the Moral Sciences* (London: Hentley, Boston: Routledge and Kegan Paul).
- (1978b), 'Realism and Reason', in *Meaning and the Moral Sciences* (London: Routledge and Kegan Paul).
- Quine, W. V. (1953), 'Two Dogmas of Empiricism', in *From a Logical Point of View* (Cambridge, Mass.: Harvard University Press), 20–46.
- (1960), *Word and Object* (Cambridge, Mass.: MIT Press).
- Strawson, P. F. (1974), *Subject and Predicate in Logic and Grammar* (London: Methuen).

# 8 Proper Function and Convention in Speech Acts\*

---

## 1. Introduction

Strawson's 'Intention and Convention in Speech Acts' (1964) introduced into speech-act theory two of its most characteristic contemporary themes. Strawson applied Grice's theory of communication to speech-act theory (Grice 1967). Using this tool, he then drew a distinction between two kinds of illocutionary act. I will prosaically call these two 'K-I (kind I) speech acts' and 'K-II (kind II) speech acts'. Strawson claimed that contrary to Austin's views, only K-II acts are 'essentially conventional'. Elsewhere I have complained against the first of Strawson's innovations, indeed against the whole of the Gricean theory of communication (Millikan 1984: ch. 3).<sup>1</sup> But it is not necessary to embrace the details of Grice's theory to appreciate the main shift of view from Austin to Strawson on K-I acts.

Austin had taken all illocutionary acts to be differentiated and defined according to conventional roles they were playing: in the absence of conventions to determine these roles, performances of these acts would be strictly impossible. Strawson claims that there is a large class of illocutionary acts, the K-I acts, that are differentiated not by reference to conventional roles of any sort, but by reference to the purpose of the speaker in speaking. Other philosophers soon took sides with Austin or Strawson. For example, Schiffer (1972), Bach and

\* This is a heavily revised version of 'Proper Function and Convention in Speech Acts' from Lewis E. Hahn, (ed.), (LaSalle, Ill.: Open Court, 1998), 25–43; © *The Philosophy of Peter F. Strawson*. The Library of Living Philosophers 1998, used with permission.

<sup>1</sup> For a good discussion of this issue see also Recanati (1987).

Harnish (1979), and Recanati (1987) agreed with Strawson that K-I acts are defined according to certain kinds of Gricean intentions expressed by speakers, while Warnock (1973) and Searle (1989) sided with Austin, claiming that the difference between K-I and K-II acts is only that the former invoke merely linguistic conventions while the latter invoke wider social conventions. All seem to agree, however, that despite some borderline cases, there is a fundamental difference between the kinds.

In this chapter I will explore certain relations between *purpose* and convention, intention being, of course, a kind of purpose. I will describe three aspects of speech acts that are characteristically tightly interlocked:

- (1) the speaker's purpose in speaking;
- (2) the purpose (function) of the public-language expression used;
- (3) the conventional extralinguistic move made, if any, classified by its conventional outcome (which typically accords with its conventional purpose).

I will argue that Austin's and Strawson's (and Grice's and Searle's, etc.) analyses all suffer from failure to recognize the very existence of purposes of the second sort—the independent functions of public-language expressions—and failure to discuss purposes of the third sort. As a result they fail to recognize that the main elements defining illocutionary acts are the purposes behind them. In the case of many of the speech acts that Austin had centrally in mind when he coined the term 'illocutionary act', more than one of these kinds of purposes is present, these typically coinciding in content, but there are also cases in which these multiple purposes diverge. The names and descriptions offered by Austin as designating specific 'illocutionary acts' cover various of these kinds of purposes. Strawson's description of K-I and K-II categories roughly corresponds to the difference between purposes of kind (1) and/or kind (2) and purposes of kind (3). Where Strawson argued for a two-way distinction despite some borderline cases, I will argue for a continuum from (1) to (3), depending on how well the contents of these three kinds of purposes overlap or fail to overlap.

To make this argument I will need to call in a model of purposiveness under which purposes can be univocally attributed to linguistic expressions, to conventions, and to speakers. I will also need to call in a different model of the conventionality of language than has generally been employed. The model of purpose that I will use is developed in detail under the label 'proper function' in my (1984: chs. 1–2; 2002: chs. 1–2). The theory of convention is detailed in Chapter 1 of this volume. Here I will sketch these models only with very broad strokes.

## 2. *Strawson's Distinction*

Strawson's original suggestion was this. K-I illocutionary acts are completed when the hearer recognizes that the speaker has a certain kind of intention in speaking. This intention is to secure a certain response from the hearer, such as an action (paradigm imperatives) or the forming of a belief (paradigm indicatives). Further, in Gricean fashion, the speaker's, S's, intention is to procure this response by means of the hearer's, H's, recognizing that S intends or wishes to procure it, and by H's recognizing that S intends H to recognize this latter intention, and so forth. For example (I am interpolating here; the examples are not Strawson's), differences among the illocutionary acts of reminding, informing, asking, and warning lie in the responses or effects the speaker intends to produce in the hearer. In the case of reminding, the intended effect might be getting H to recall, for informing getting H to believe, for asking getting H to impart certain information to S, and for warning bracing H against dangers that S describes. Various K-I acts that involve intending identical hearer responses are then further differentiated in accordance with more exact mechanisms by which the speaker intends or expects to procure the intended response. For example, in the case of requests versus entreaties (these are Strawson's examples), S intends H to understand how he holds the intention that H comply, whether 'passionately or lightly, confidently or desperately' (1964: 610), and S intends that this knowledge should motivate H to comply; in the case of orders, S intends that



H understand that the context of utterance taken together with certain social conventions implies that certain consequences may follow if H does not comply, intends that this knowledge should motivate H, and so forth.

K-I acts, Strawson says, are not 'essentially conventional'. It is not true that acts of warning, acts of entreaty, acts of requesting, and so forth

can be performed only as conforming to . . . conventions . . . to suppose that there is always and necessarily a convention conformed to would be like supposing that there could be no love affairs which did not proceed on lines laid down in the *Roman de la Rose* or that every dispute between men must follow the pattern specified in Touchstone's speech about the countercheck quarrelsome and the lie direct (p. 603)

Nor, when a speaker disagrees with someone, is there, in general, some '*convention* that constitutes' his act as an act of raising objections. So far so good. This may allow us to understand how various illocutionary acts can be performed without using 'explicit performatives' such as 'I insist' or 'I warn you' or 'I promise'.

On the other hand, Strawson tells us, the speaker may explicitly avow his illocutionary intention, thus conveying the force of his utterance conventionally, saying, for example, 'I warn you that . . .' or 'I entreat you to . . .' or 'I object to . ..'. In the case of these 'explicit performatives', the 'conventional meaning' of the expression used may 'completely exhaust the illocutionary force' of the utterance.

As examples of K-II illocutionary acts, Strawson lists 'an umpire giving a batsman out, a jury bringing in a verdict of guilty, a judge pronouncing sentence, a player redoubling at bridge, a priest or civil officer pronouncing a couple man and wife' (p. 611). Here, Strawson says, the intention of the speaker is not to secure a particular response from the audience. Such acts are performed according to the rules of certain conventional procedures (for example, the marriage ceremony) as acts 'required or permitted by those rules . . . done as falling under the rules . . . the speaker's utterance is not only *intended* to further, or affect the course of, the practice in question in a certain conventional way; in the absence of any breach of the conventional

conditions for furthering the procedure in this way, it cannot fail to do so' (p. 612).

Unlike K-I acts, K-II acts do not, as such, have as ingredients speaker intentions whose fulfillment is dependent on hearer cooperation. Reciprocally, unlike K-II acts, 'the wholly overt intention which lies at the core of . . . [K-I acts] . . . may, *without any breach of rules or conventions*, be frustrated' (p. 613; emphasis in the original). What is common to the two kinds of illocutionary act is that the speaker's intention is 'wholly overt'; that is, intended to be recognized by the audience. What is different is that in one case the intention is to produce a certain audience response, in the other to 'further a certain [conventional] practice' (p. 612).

Strawson concludes with the caution, 'it would certainly be wrong to suppose that all cases fall clearly and neatly into one or another of these two classes . . . [for example a] speaker whose job it is to do so may offer information, instructions, or even advice, and yet be overtly indifferent as to whether or not his information is accepted, his instructions followed, or his advice taken' (p. 614).

### 3. *The Disunity of K-I Categories*

Despite some ambiguity in the passage just quoted, it is clear that Strawson intends the K-I/K-II distinction to apply not to instances but to categories of acts, designated, for example, by performative verbs such as 'inform', 'instruct', 'advise', 'entreat', and so forth. Consider, in this light, those explicit performative K-I instances (such as instances containing 'I warn you . . .', 'I advise you to . . .') about which he says that the 'conventional meaning' 'exhausts the illocutionary force'. What is the relation supposed to be here between convention and force?

Strawson cannot mean here that these explicit instances are conventional in the sense that their having the force they do is constituted by convention, that there are conventions that make them into acts, say, of advising, warning, entreating, objecting, and so forth. A class of speech acts grouped together and named according to the particular

audience response intended by the speaker cannot be, as such, a class that an instance gets into by convention. That a speaker has a given intention cannot be a mere matter of convention. There cannot be a convention that turns something else into a speaker intention. There might of course be a convention always to *treat* certain kinds of actions, for certain purposes, *as though* they embodied certain intentions—to 'count them as' embodying these intentions; that is, as legally or morally or socially equivalent to actions actually embodying these intentions—but this would not make them into embodiments of intentions.

Rather than an explicit performative turning a K-I act into an act *constituted* by convention, perhaps Strawson means that explicit K-I acts are natural acts of expressing one's intentions that happen to be performed in a conventional manner. For example, one can perform the natural act of holding one's fork in a conventional manner; say, tines up, in the right hand, with the thumb on top. One expresses one's intention to warn by saying 'I warn you' or expresses one's intention to advise by saying 'I advise you'. But, of course, saying one has a certain intention does not make it so. Perhaps this is why Strawson cautions that he is considering in his essay only 'normal and serious' uses of language (p. 599), which might mean here that we are to consider explicit performatives only when backed by the intentions these conventionally express. Thus, 'I warn you that she is exceedingly charming' probably would not express a speaker's intention to brace the hearer against a danger, but neither would it be 'normally and seriously' used. This would make it possible to argue, as Strawson does, that by the use of explicit performatives of the K-I kind nothing is accomplished that could not in principle have been accomplished in a world without conventions. It is merely that something natural has been done in a conventional *way*. It is not really the speaking of certain words, then, but merely the expressing of a certain audience-directed intention that makes an utterance of, say, 'I entreat you . . .' into an entreaty, and so forth.

The trouble with this suggestion is that besides uses of explicit K-I performatives that are not literal or serious and not intended to be understood as literal or serious there are also K-I acts such as warning,

requesting, and entreating that are performed with the intention that they should be taken seriously yet without the speaker actually intending the outcome that seems to define these acts—a lie is not just a nonliteral usage. For example, it is perfectly possible for a person to use an explicit performative in order to request, entreat, order, or demand a thing that she does not intend the hearer to accomplish, intending only to distract the hearer, or to embarrass the hearer, or to trick the hearer into failure or into starting into the designated action. Nor can we patch up this suggestion by requiring for a K-I illocutionary act only that S intend H to think that S intends H to act or believe. Strawson himself gave us the contrary example: 'A speaker . . . may . . . be overtly indifferent as to whether or not his information is accepted, his instructions followed, or his advice taken.' Faced with this sort of difficulty, some authors (Schiffer 1972, Bach and Harnish 1979, Recanati 1987) have proposed another epicycle: the K-I act requires only that S intends that H recognize S's intention to 'provide H with reason to believe' that S intends H to act or believe. The difficulty of understanding this claim, and the fear of provoking yet more epicycles, discourages discussion. A less baroque solution seems preferable.

I will propose that K-I acts are defined quite straightforwardly by their *purposes*, and that the difficulties with regard to them stem from the fact that there are two different kinds of purposes that may be involved. Besides the purposes or intentions of the individuals who use various conventional language forms there are the purposes of these very forms themselves. Because conventional language forms have their own purposes—their own functions—their instances can fall into K-I categories directly, without accompanying speaker intentions. Difficulties in classification of certain speech-act instances arise because these two kinds of purposes sometimes conflict.

To explain this I will first need to explain what conventions are in the broad sense of 'convention' that covers linguistic conventions. Then I will talk about purposes and how conventions acquire them. This will reveal how a conventional illocutionary act can have its own purpose, additional to the purpose of the speaker. Last I will discuss a still broader sense of 'convention' which covers the conventionality of K-II illocutionary acts and I will discuss their purposes.

#### 4. Conventions

It is the conventionality of *acts* and *activities* and *patterns of activity* that will concern us—as a compromise, often I will just say ‘activities’. In the sense of ‘conventional’ that we need to examine first, conventional activities are ‘reproduced’ items. Instances of conventional activities are modeled on prior instances of that activity, previously performed, typically, by other people. More formally, the concrete form—I am going to say ‘shape’—of the natural activity embodying an instance of a conventional activity is determined according to the shape of certain historically prior natural activities such that had the shape of these prior activities been different along certain dimensions, this instance would have differed accordingly. This may be because the activity is directly copied or imitated. Alternatively, the process of reproduction may be indirect, as when one person instructs or trains another to act in accordance with a convention.<sup>2</sup> Second, when they have functions, as they often do, conventional acts, activities, and patterns of activity are characterized by a certain arbitrariness in relation to function. Thus, patterns of skill, though they may be handed down by copying and instruction, are not, in general, conventional patterns. A conventional activity is one whose conventional shape is not wholly dictated by its function. A reproduced activity or pattern that exemplifies a skill is conventional only if it might have differed or been replaced by differently shaped activities or patterns which, assuming similar proliferation in the culture, would then have produced the same result. This does not give us a sharp boundary for the conventional. A borderline case might be, for example, certain ‘conventional’ techniques handed down in different schools of violin playing. These techniques are not totally interchangeable. They have subtly different musical effects. But the conventions of each school taken *in toto* accomplish pretty much the same.

An instance of conventional activity is such in part because it has been reproduced. That is, certain aspects of its shape have been

<sup>2</sup> For a more formal discussion of ‘reproduction’ as that term is meant here see my (1984: ch. 1).

reproduced. Other aspects of the instance will not have been reproduced, of course—details of how it is performed, where, by whom, and so forth. The instance is conventional, then, under a description of its conventional shape, pattern, or place. *Wearing black to a funeral* describes the shape of a certain convention; *shaking hands with the right hand* describes the shape of another. But *wearing black to a funeral* and *shaking hands with the right hand* are not names of shapes that are conventional merely *as such*. Instances falling under these descriptions are conventional only when done *as* following convention; that is, when reproduced rather than accidentally instantiated. Thus,

(1) It is conventional to wear black to a funeral.

and

(2) Susan wore black to a funeral.

when conjoined, do not imply

(3) Susan performed a conventional act.

Compatibly, there is nothing to prevent the same shape from being used in more than one convention. Raising one's hand is conventional in order to vote. It is also conventional in order to request to speak. Which if either of these conventions is instanced when a particular hand is raised depends on which if any previous instances of hand-raisings are the causes and models for this one, instances used for voting, or instances used for requesting to speak.

Some conventional patterns of activity are very complex. The conventional pattern that is reconstructed when a group of children plays *ring-a-ring o' roses* is rather complex. More complex are the conventional patterns that are reproduced when parliamentary procedure is followed during a meeting, or when one plays chess. These latter patterns are not only complex but 'relational'. It is not absolute shapes (forms) that are reproduced but relationships between shapes. One can follow parliamentary procedure in any language while discussing issues of any content, and chess can be played with any shaped pieces; for example, with bottles on a sandy beach. Such patterns are most easily described by giving 'rules' for their construction, often rules involving conditionals. These are rules that *have* to be followed insofar—and *only insofar*—as one's purpose is to follow the particular

convention. For example, you must follow the rules of chess if you wish to play a conventional game of chess; otherwise you can do what you like with the pieces.

Many of the patterns of activity that define, within a given society, the institutions of marriage, property holding and transfer, use and transfer of conventional powers, and so forth are complex, relational conventional patterns. Many may also be written into law or other sanctioned regulations. Others are *only* written into law or other regulations. They are never or almost never merely reproduced from previous examples of the same pattern. Then, they are not conventional in the sense I am now discussing (but see Sect. 6 below).

Reproductively established conventional patterns are often reinforced with sanctions of one kind or another. This has no bearing on their conventionality. Some cultures frown on every failure to conform to convention. But it is neither the threat of these frowns nor the degree of conformity to a convention that makes it into a convention. A substantial literature to the contrary, conventions are not, *as such*, either shapes that everybody in some group conforms to, or shapes to which everybody in some group thinks you are supposed to conform. Think, for example, of wearing white for tennis, which in many circles is conventional even though nobody much cares whether you do or you don't. And think about moves in chess. After any given move in chess one could always quit and start playing dolls with the pieces. If that is not allowed, it is rules of etiquette or tournament rules, not chess rules, that prohibit it. The rules of chess don't tell you that you can't quit, but only what would constitute going on.

Nor should we suppose that conventions are instantiated only by people knowingly following them. Witness the conventions for correct social distance when conversing. These distances vary from culture to culture, and are unconsciously reproduced by being learned as a skill. If you are at the wrong social distance, the one to whom you are speaking will move, so that to avoid slow circling about as you talk, you learn to stay at the conventional distance. Similarly, a person, even everyone, might unconsciously learn to conform to the convention of driving on a given side of the road solely as a skill—as a means of avoiding oncoming traffic.

One particularly interesting kind of piece of a conventional pattern deserves a name of its own. This is an optional piece of a conventional pattern, the shape of which piece puts constraints on the shape of ensuing pieces. That is, if the conventional pattern is to continue to unfold, the rest must be conformed to this piece, so as to bear the right relation to it. Thus, a conventional move during a chess game and a conventional Protestant marriage ceremony each constrains what can follow while according with relevant conventions. Such acts are interesting because convention does not say when to perform them, yet they effect changes in the situation that must subsequently be taken into account if events are to unfold under the covering conventions. They are 'free' actions under the convention that have predictable effects under the convention or that put constraints on what can happen later under the convention. Thus, they put constraints on *actual* effects insofar as the convention is *actually* followed. We can call such actions 'conventional moves' and speak of them as having 'conventional outcomes'.

Contrary to the flavor of recent discussions, probably I should explicitly mention that the situation with conventional moves is *not* this: that having instanced a certain action shape under certain (difficult-to-pin-down) conditions automatically 'counts as' x-ing, irrevocably, inexorably, no matter how much you kick and scream. In particular, there can be no such thing as a conventional 'rule of the form "X counts as Y in context C"' (Searle 1969: 52) where 'X' is the description merely of an activity's shape. First, instancing an action shape is never, just as such, instancing a conventional move. The shape must be a reproduction; it must non-accidentally follow a conventional pattern. Second, *simply as such*, making a conventional move is merely putting in place a piece of a reproduced pattern which others (or oneself) may or may not then be motivated to complete. The pattern allows or requires a decision to be made at the location of the conventional move. In most cases participants will not be interested in following through to the conventional outcome of an apparent move if they think it was not made intentionally.

On the other hand, especially in cases where abiding by the convention happens to have sanctions attached, and especially where



there might be reason for a person sometimes to be dishonest about whether a move was intended, then, as Strawson put it, 'the play is strict' and either some wider convention or some law or regulation may override the need for a consonant intention behind a conventional shape introduced at a decision location. In such circumstances one may also expect that the wider convention, law, or regulation is quite strict about the exact shapes to be used for the moves. Sloppy reproductions will not be allowed. But this rigidity is not in the nature of conventions themselves but in the nature of the more encompassing conventions or sanctioned regulations that may sometimes apply to them. And, of course, the same is true of any more encompassing conventions involved as well. For example, the convention that any piece of your own that you touch while playing chess you then must move whether the touch was intentional or not does not enforce itself. When enforced, it is enforced by tournament regulations or by ordinary social sanctions.

Some labels group activities together merely by shape, for example *wearing black*, *wearing black to a funeral*. But there is also a way of describing conventional activities, moves, and outcomes that exploits the fact that different shapes occurring in different conventions, in different traditions, may hold corresponding places in their respective conventional patterns, hence may be classed together. 'Getting married', 'giving a (conventional) greeting', and 'observing table manners' are such labels.<sup>3</sup> Each of these role-described classes of conventional moves, activities, or patterns can take any of numerous shapes.

There may be a certain ambiguity in some of these abstract labels. For example, does one who wears black to a funeral without knowing of the convention 'wear funeral attire'? Does the cat that bats my queen next to your king 'put your king in check'? Let's just decide in the negative and speak of activities falling into 'role-described' categories only when

<sup>3</sup> Moves such as castling in chess are a little different. For although pieces of any shape can be used for chess, indeed chess can be played even without a board (for example, games can be played by postcard), games of chess all fall in the same tradition. The game of chess is an abstract pattern of singular origin that is sometimes reproduced in highly imaginative ways. (Wilfrid Sellars once suggested using Cadillacs for kings and Volkswagens for pawns, etc., and moving them from one Texas county to another.)

done *following* convention. Thus, role-described acts are acts that could not be performed were there not covering conventions. Similarly, if I speak of doing something 'in a conventional manner' or of 'making a conventional move' I will imply that it is done not accidentally but following a convention. Dressing in a conventional manner is a role-described act which takes different shapes in different cultures; the conventional move that is getting married is another role-described act which takes different shapes in different cultures.

So far as I can make out, terms naming role-described categories of conventional moves always classify these moves by conventional outcome. For example, making a bid, getting married, performing a naming ceremony, adjourning a meeting, putting a king in check, scoring a goal in sports are each role-described moves classified according to their conventional outcomes. Now, apparently Austin thought that names for K-I acts such as 'warning', 'entreating', 'ordering', 'objecting', and so forth were like names for moves in a game, in which case their instances would be classified by conventional outcomes—by what the conventional results would be. Strawson on the other hand claimed that these labels classify primarily by speaker intention. These positions are obviously incompatible: a speaker intention is not the same thing as a conventional move. My suggestion will be that K-I illocutionary acts are actually classified by purpose, the conventional outcomes of utterances of K-I explicit performatives being equivalent to the purposes of these *as public language forms*—equivalent, that is, to their linguistic functions. Classification by purpose is another way, other than classification by conventional outcome, to classify acts more abstractly than by shape. For example, what constitutes hunting is a matter of purpose, hence quite differently shaped activities can constitute 'hunting mice' as a hawk does it, as a cat does it, and as I once did it thumbing through the yellow pages in preparation for a mouse-loving daughter's Christmas. Since speaker intentions are also purposes, Austin and Strawson were both right, each in his own way.

But I have been getting ahead of my story. First, I must clarify and defend the claim that uses of public-language forms are indeed conventional moves having conventional outcomes. Then, I must

clarify and defend the claim that these outcomes correspond to functions or purposes of the language forms themselves.

### 5. *Speaking as Making Conventional Moves*

It is not, of course, a matter of convention whether a speaker has communicated his intention. Nor is it always a matter of following convention that a hearer should respond as a speaker wishes her to. But a hearer does follow convention when she does what a speaker explicitly *says* to do, or believes what a speaker explicitly *says* is true, and does so, in the normal way, *because* of what the speaker has said. In each such case the hearer completes the reproduction of a conventional pattern of movement from speaker words into hearer reactions. Correlatively, in each of these cases the speaker makes a conventional *move* having a conventional outcome. He lays down the beginning of a conventional pattern in a way that constrains what can follow *in accord with the convention*. To see that this is so it is important to keep in mind (1) that to follow a convention is not mandatory *as such*—for example, no sanctions need constrain the hearer to respond in a conventional way—and (2) that following conventions is not always following conscious rules. The speaker's production of the expression and the hearer's cooperative response to it constitute a reproduced pattern whose form is arbitrary relative to its coordinating function. That is all that is needed for convention. Contrast the conventional syntactic and tonal patterns that embody tellings to and tellings that with conventional exclamations ('Hurrah!', 'Ouch!') which are merely conventional *means of expression*, calling conventionally for no particular determinate response from the hearer.

In the case of conventional directive uses of language such as paradigm uses of the imperative mood, the pattern that is conventionally reproduced begins with an intention or desire of S's that H should act in a certain way. It is completed when H has acted that way as a result of guidance, in accord with conventional rules for guidance, from conventional signs made by S. That the pattern is not completed until H has acted as directed is clear, for new instances of the pattern would not

be initiated by speakers were it not that hearers sometimes complete such patterns. The first part of the pattern is conventional, is reproduced, only because both parts are sometimes reproduced. Thus, when you ask or tell me to do something in a conventional way, using some appropriate shape from some public language to do so, it is *conventional* for me to comply: this outcome is the completion of a conventional pattern. The full recipe for the convention tells what the hearer has to do on hearing such-and-such words.

Similarly, when S tells H that something is the case in a conventional way, it is conventional for H to believe it. That the pattern is not completed until H has been guided into belief in accordance with the conventional rules is clear because new instances of the pattern would not continue to be initiated by speakers were it not that hearers sometimes believe what they are told. Had earlier hearers responded in accordance with different patterns of interpretation and belief this would have affected the behavior of speakers, with the result that H would have learned to exhibit different patterns as well. These patterns of belief-formation do not result from voluntarily following a rule, of course, any more than standing at the right social distance is a result of voluntarily following a rule. Whether or not I believe what I hear is not under voluntary control. But believing what I hear results, in part, from a process of reproduction enabled by learning. I believe as I do in response to what I hear in part because others who speak the same language have followed similar patterns in moving from what they have heard to what they believed, thus reinforcing speaker uses of these language forms. I come to believe what I hear as conforming to a convention.

Consider now what the speaker–hearer patterns will be like that characterize the uses of explicit performatives mentioning K-I acts. Suppose that Strawson is right: For S to convey certain intentions as to H's response can be a way (if not, I shall soon argue, the only way) of performing a K-I act such as requesting, warning, informing, and so forth. If the full pattern that is conventional when S tells H that something is the case is as described above, then S's explicitly telling H that S is performing a certain K-I act has as a conventional outcome that H believes S is performing that act. That is, the conventional

outcome will be that H believes that S has the intention associated with that K-I act. And insofar as Hs are sometimes obliging, they will sometimes fulfill the speakers K-I act intentions. (If Hs were never obliging, speakers would soon stop having K-I intentions.) A repeated pattern will thus emerge that begins with a speaker's K-I intentions about a hearer's responses, moves through explicit mention of a corresponding K-I act ('I warn you that . . .', 'I entreat you to . . .'), and ends with the hearer fulfilling the speaker's intentions. This successful pattern will begin to be *reproduced*, these uses of explicit performative K-I expressions will soon becoming fully conventionalized. Thus, the use of an explicit K-I performative will acquire the accomplishment of the K-I act it designates as a conventional outcome. But then Strawson was wrong on one point. He was wrong that when explicit performatives are used to perform K-I acts 'the wholly overt intention which lies at the core of [K-I acts] . . . may, without any breach of rules or conventions, be frustrated' (p. 613, emphasis in the original)—assuming, that is, that the speaker's intention accords with the performative verb she uses.

We have seen that speaking consists in large part of making conventional moves. In particular, purposefully to use an explicit K-I performative verb in the first person is to perform a conventional move. The move made—a move classified, as usual, according to its conventional outcome—is named by the performative verb that is used. The speaker who, following the standard convention, says 'I warn you that . . .' performs an act of warning whether or not the particular hearer follows through conventionally by heeding the warning, and the speaker who, following the standard convention, says 'I entreat you to . . .' performs an act of entreating whether or not the particular hearer follows through conventionally by complying. Just as certain forms intentionally having been gone through entails that you are married, or have been christened, or that the meeting is adjourned—these are descriptions of conventional moves classified by conventional outcome—so it is that certain verbal forms having been reproduced intentionally make it the case that you have been warned or entreated or informed and so forth.

Thus, it may appear that Austin was right about the conventional nature of illocutionary acts using explicit K-I performative verbs while

Strawson was right about those K-I acts that are not performed by explicitly naming themselves, acts where the hearer is expected to gather in some other way what the speaker's fine-grained illocutionary intentions are. On the other hand, surely the typical use of an explicit performative K-I verb will be by a speaker who intends its conventional outcome. So both Austin's analysis and Strawson's seem to fit these common cases. Strawson's analysis might be considered the more general analysis, failing only in cases of insincerity or, as he himself suggests, in those peculiar cases where a speaker merely goes through the correct form of using an explicit performative verb of warning, informing, or whatever while being 'overtly indifferent as to whether or not his information is accepted, his instructions followed, or his advice taken'. Still, the existence of these cases make it look as though K-I acts are not really all of one kind. There seems not to be anything common to them all. How does the mere fact, in these latter peculiar cases, that a certain conventional move has been made get into the same category as the fact, on other occasions, that a speaker has a certain intention in speaking?

This issue can be resolved by a look at the purposes (functions) of public-language forms. If the conventional outcome of use of an explicit performative can be shown to be the same as its purpose qua conventional language form, then K-I illocutionary acts are all classified in the same way after all—classified by their purposes. The exceptions will merely be cases in which although the language form itself has the relevant purpose, the speaker does not.

## 6. *K-I Acts are Defined by their Cooperative Proper Functions*

We say of cats and dogs that they 'ask to come in' or 'beg to go for a walk'. It is quipped that a difference between dogs and cats is that 'dogs request while cats demand'. There are interesting studies said to be about animals' 'greetings to conspecifics', 'invitations to play', 'warning cries', and so forth. Assuming that animals do not indulge in embedded

Gricean intentions, these uses of illocutionary verbs do not fit Strawson's account of K-I acts any better than those cases of informing, warning, and so forth mentioned above that are performed despite the speaker's insincerity or overt indifference about the hearer's response. I wish to propose a simpler account than Strawson's according to which each of these descriptions is literally correct. K-I acts are grouped under their appropriate labels by, in a suitably broad sense, their *purposes*—more precisely, in an idiom which I will explain in a moment, by their 'cooperative proper functions'.

Here I can offer only a quick review, lacking defense, of the notion 'proper function'—just enough to give the flavor.<sup>4</sup> Items have proper functions as belonging to families of items reproduced one from another, where the continued reproducing depends or has depended on some function these items serve. The biological functions of body organs and the functions of mechanisms that produce tropistic behaviors are prime examples of proper functions. The functions of behaviors learned by trial and error are also examples, as are the functions of customs (when customs have functions) and the functions of words and syntactic forms. An item also has a proper function if it is the product of a prior device designed to vary or adapt its productions depending on circumstance so as to perform certain functions in those circumstances. These adapted productions then have 'derived proper functions'. Using a well-worn example (well-worn by me), the mechanisms in worker honey bees that produce bee dances are supposed to vary their dances depending on where nectar is located so that the dances can guide fellow bees accordingly. The different dances that result have different derived proper functions: each is supposed to send watching bees off in a different direction. I have argued that behaviors produced in a normal way by the behavior-producing mechanisms in humans and other higher animals also have derived proper functions, though usually derived in a far more complex way, and that these functions coincide with what we would usually identify as the purposes of these behaviors or of the individuals exhibiting them. Human intentions, understood as

<sup>4</sup> The notion 'proper function' is defined in my (1984 chs. 1, 2). It is further explicated and applied in my (1993; (e.g.), chs. 1, 2, 11) and in my (1994; 2002).

goal representations harbored within, have as derived proper functions to help produce their own fulfillments.<sup>5</sup> And human artifacts have as proper functions the purposes for which they were designed.

Now consider the syntactic forms for indicative, imperative, and interrogative moods in the various natural languages.<sup>6</sup> These are reproduced from one speaker to another; children copy them from adults. These forms continue to be differentially reproduced because they are serving differentiated functions. And they serve differentiated functions because hearers respond to them differentially. What stabilizes these functions?

The evolutionary mechanism at work here is parallel to that which tailors the species-specific song of a bird and the built-in response of its conspecifics to fit one another, or tailors the nipple of the mother and the mouth of her infant to fit one another, but with learning standing in for natural selection. Speakers (collectively) learn how to speak and hearers learn how to respond in ways that serve purposes for them, each leaning on the settled dispositions of the others. This kind of co-tailoring requires there to be functions served at least some of the time through cooperation between the partners—enough of the time to keep their cognitive systems tuned to one another. So there must be purposes that are sometimes served for hearers as well as for speakers, served by hearers' responding in a cooperative way to what speakers say. For example, often enough there are rewards to motivate hearers who satisfy imperatives. Simply pleasing the speaker may be enough reward, but various kinds of sanctions may also apply. And speakers often enough speak the truth, so that there are often rewards for hearers who believe indicatives. Indeed, if speakers failed to speak the truth too much of the time, there would be no way for hearers to learn a language.<sup>7</sup> Thus, a proper function of the imperative mood is to induce the action described, and a proper function of the indicative mood is to induce belief in the proposition expressed.<sup>8</sup>

<sup>5</sup> See esp. my (1984: ch. 6; 1993: ch. 8 sect. 6; 2004: chs. 1, 5, and 6).

<sup>6</sup> For more details see my (1984: introd., chs. 1–4), and Ch. 3 to this volume.

<sup>7</sup> See Ch. 10 to this volume.

<sup>8</sup> See my (1984: ch. 3). It is important that these proper functions are not derived merely by *averaging* over speaker intentions or hearer reactions.



It seems, then, that the proper functions of the grammatical moods accord with the conventional outcomes of the conventional moves made in using them. Turning to the functions of explicit K-I performatives, the same principle seems to hold. For example, the survival value of the expression 'I warn you that p' would seem to be that it sometimes causes hearers to be braced against dangers connected with p, and the survival value of 'I order you to A' would seem to be that it sometimes causes hearers to A so as to avoid negative sanctions, and so forth. A proper function of an explicit K-I performative is to induce the hearer response that is the conventional outcome of the conventional move made in uttering it.<sup>9</sup>

So we have come full circle. Austin was right that using an explicit K-I performative is making a conventional move; conventional moves are classified by their conventional outcomes; and a proper function or linguistic purpose of a K-I performative language form is to produce its conventional outcome. Strawson was right that K-I acts performed without the use of explicit performative are classified by speaker intention, by the purpose of the speaker in speaking. It follows that all K-I acts, performed with or without the use of explicit performative verbs, are classified by their purposes. Some are K-I acts because of the speaker's purpose, some because of the purpose of the language form used, and most, perhaps, for both reasons, since these two purposes are likely to match.

This result can be generalized. If a K-I act is defined primarily by its purpose, then the dog asking to go out at the door does indeed perform a K-I act of requesting, as does the person gesturing in a nonconventional but polite manner for you to open the door, the person who says 'Please open the door', and the person who says 'I request that you open the door'. All are behaving in ways that have as their purpose to get you to open the door (and without special sanctions attached). Also, in each of these cases the proper function of the act is, more specifically, a 'cooperative proper function'. If fulfilled in the normal

<sup>9</sup> Not every conventional move has a proper function. Conventional patterns may sometimes be imitated for no particular reason rather than due to some function they are performing. Functional convention blurs into the pointlessly convention-bound and into mere habit.

way<sup>10</sup> it will be fulfilled through a cooperating response; that is, a response that has as part of its own proper function to complete the initiating act's proper function. Note the close analogy here to the reciprocal structure of Gricean intentions. According to Grice, the hearer is to do as the speaker intends in part because the speaker intends it.

The proposal then is that K-I acts are defined by their cooperative proper functions. Because the various grammatical moods have cooperative proper functions, their use *in and of itself* is enough to constitute a broad kind of speech act—at least, for example, a *telling that* or a *telling to* or an *asking whether*. More specific K-I speech acts not employing explicit performatives are then differentiated according to additional or more fine-grained *speaker* (rather than expression) purposes—more exact mechanisms by which the speaker intends or expects to procure the intended response—just as Strawson said. In the case of the explicit performatives, K-I acts are differentiated by naming themselves. Thus, they become 'self-verifying'. As instances of the original informative conventional pattern from which they were derived (Sect. 5 above) and which they also exemplify (they remain in the indicative mood), their more recent conventional purposes (linguistic functions) accord also with what they *say* is the purpose of the saying. For example, 'I warn you that *p*' has both the function of informing you that it is a warning and also the function of warning you, whether or not you follow through with the conventional responses to this form, hence whether or not its conventional purpose is fulfilled.

It remains to discuss cases in which the intention of the speaker does not accord with the proper function of the language form used. Strawson gave us examples in which the speaker is overtly indifferent about whether his speech act actually performs its proper function, but perhaps cases of actual conflict are more interesting. For example, the person who says 'Open the door!' may be sarcastic or joking or acting in a skit; she may intend only that you trigger a booby trap placed over the door; she may misunderstand the public-language function of the words she uses. If she actually intends something other than your opening the door, then her saying 'Open the door!' has two conflicting proper

<sup>10</sup> If fulfilled, that is, in accordance with a 'normal explanation'. See my (1984: chs. 1, 2).

functions at once. The first is derived from the history of the language forms she reproduces (the imperative mood and the words she arranges into this mood). This is the proper function of the *expression* she uses. The second is derived from her intention in speaking.

In the case of insincere uses (the booby-trap case) the proper function derived from the speaker's intention is not a *cooperative* proper function, so it does not affect the question what sort of K-I act is being performed. Suppose that Strawson is right that K-I acts are defined by their *cooperative* functions, not merely by the speaker's intention in use, rather as Grice proposed for properly communicative uses of language. In insincere cases, then, the only speech act performed is the act defined by the purpose of the language form itself. If she says 'Open the door!' intending only that you try to open it and thus trigger the booby trap overhead, she directs or tells you to open the door even though that is not at all what she intends you to accomplish.

But there are also cases where a speech act has conflicting functions both of which are cooperative. For example, if I jokingly ask you to leave I do not intend you to leave but I also do not intend you to react as if I were really asking you to leave. I intend you to understand that I am only pretending to ask you to leave; my purpose is cooperative. Or consider these cases: the armed robber smiles and says 'I entreat you to hand over your money'; Anytus threatens Socrates, 'I advise you to be careful'; Mom orders 'I am asking you for the last time whether you are going to take out the garbage!'; the friend who has matched Jim with a blind date says 'I warn you that she is exceedingly charming'; Father warns 'If you aren't back by twelve I promise that you will regret it'. The proper functions of these expressions do not match the cooperative intentions of the speakers. Has the designated act been performed? These may be uses that Strawson intended to rule out as not 'literal and serious' uses. Certainly, these are not paradigm cases of asking, advising, entreating, warning, or promising. It does seem natural to say, however, that the act performed is the one according with the cooperative speaker intention rather than with the purpose of the linguistic form used. Anytus has threatened Socrates, Father has given a warning, and so forth. These are acts in which the speaker intends the hearer to understand his actual intentions despite the ironical language used. Here, in attributing a speech act, the 'wholly overt intention' of the speaker to do something

other than what he explicitly says he is doing seems to override the purpose of the language form he uses.

Yet things are not quite that simple, for it is possible to question Strawson's claim that a 'wholly overt intention . . . lies at the core' of all K-I acts even in certain cases that cannot be explained away by pointing instead to the purpose of the language forms used. Sperber and Wilson have questioned whether, as a general rule, speech acts 'have to be communicated and identified as such in order to be performed' (1986: 244). They suggest that, for example, predicting, asserting, hypothesizing, suggesting, claiming, denying, entreating, demanding, warning, and threatening do not (p. 245). If this is right, then some of the differentia separating K-I acts performed without the use of explicitly differentiating performative verbs may not correspond to cooperative purposes of the speaker. For example, perhaps I can successfully warn you by making you alert to a danger without your understanding that as the purpose of my remarks. And if this is the case speech acts such as boasting should also count as illocutionary acts even though no such form as 'I boast that . . .' is possible—this for the obvious reason that the hearer is not going to cooperate in being impressed because he knows the speaker intends it. Suppose that a mountain climber says 'I admit that it was absolutely excruciating there towards the end', his intention being not to own up to a weakness but rather to inspire admiration for his fortitude. Despite the explicit 'I admit', this does not seem to be an admission; it seems to be a boast. Boasting is not a cooperative act, and yet the purpose of the speaker seems to override the purpose of the language form he uses in determining what speech act is considered to have been performed. Perhaps this is because the act of boasting *normally* occurs without the speaker intending cooperation, but only a certain hearer reaction.

## 7. *K-II Acts*

Not all conventional moves made in speaking are conventional moves in the sense defined in Section 4 above. Some are moves in patterns that are not reproduced from prior instances but dictated by law or

other explicit regulation. The pattern of moves required to make a foreign-born person into a US citizen, including the necessary taking of oaths and so forth, is an example of such a pattern. Also, many conventional moves are of intermediate status. The patterns in which they are embedded are at least partly copied but they or portions of them are also written into codes or laws. Marriage ceremonies, including the act of signing certain documents, are an example of this. There is no sharp line, then, between moves that are conventional in the sense defined in Section 4 and moves that are conventional because they fall under explicit regulation. In speaking of moves at or close to the explicitly regulated end of the spectrum we can speak of 'regulated conventions', of 'regulated moves', and of 'regulated outcomes'. The regulated outcome of a move may also be its proper function in a sense that falls quite strictly under the definition of that term given in my (1984). This purpose will be derived from the intentions of the person or group responsible for defining the move and its outcome—those responsible for initiating the regulations. Thus, 'in some cases, [the speaker] may be seen as the mouthpiece, merely, of another agency' (Strawson 1964: 614).<sup>11</sup>

Regulated moves, like conventional moves, are classed, as such, by their conventional outcomes. I have argued that K-I illocutionary verbs classify by purposes and that these purposes accord with the conventional outcomes of the moves made by these expressions when used performatively. K-II verbs often classify acts by conventional outcome alone. Thus, these designate conventional moves as such. These are moves whose outcomes, for any of a number of reasons, could not reasonably be intended by speakers in the absence of extralinguistic social conventions concerning them, and, indeed, whose outcomes may routinely fail to be strongly intended by speakers. Consider, for example, paradigm K-II acts such as pronouncing a couple man and wife or formally granting someone a degree. These have conventional outcomes that involve the coordination of multiple persons, many of whom may not be present at the ceremony,

<sup>11</sup> Once again, however, we should not assume that behind the conventional outcome of every regulated move there lies a clear purpose.

in behaviors forming a complex interwoven pattern, difficult even to specify. It is difficult to sum up, for example, exactly what you must do in order for you to behave in the conventional way as one who is married and what others must do in order that you be treated by them and by the law in that way. On the other hand, all or many of the conventional outcome behaviors that define K-II acts may be heavily sanctioned as falling under the covering regulations or conventions, so that they are strongly motivated. For this reason, the speaker is likely to *expect* follow-through even if he does not positively intend it.

It is generally true that the shape of a K-II move includes not just some words spoken but also the position or office of the speaker and the context of the speaking. You can't make the conventional move of bidding two no trump if it's not bridge or if it's not your turn to play, or adjourn the meeting if you are not in a meeting or not chair of it. Such requirements were labeled 'felicity conditions' by Austin, but lacking them is in fact lacking part of the move's very shape. What is reproduced or regulated as part of a pattern is not words but words-in-a-context. Hence of course the speaker cannot hope to effect the conventional outcome of the pattern out of context merely by conveying his intentions to effect it. Unlike a K-I act the K-II act can only be performed conventionally. On the other hand, it may be considered, on occasion, that the very shape of a move is, just, the conveying of the intention to perform the move by the right person at the right time. For example, at the right time in an informal bridge game one can pass with an understandable gesture, and I once witnessed a marriage ceremony performed by a severely handicapped minister who administered the vows and blessings without words. In both cases, however, the conventional setting is or was essential.

It is true that a K-II act speaker is likely merely to expect, rather than strongly to intend, much of the conventional outcome of his K-II act—as a chess player expects, rather than strongly intending, that his opponent's future moves will be constrained by his own move in accordance with the rules. Indeed, the K-II speaker may have no personal interest whatever in the outcome of his move, which he merely performs in line with custom or duty. The minister may merely be doing what he is asked to do in performing the marriage

ceremony; the provost is obliged formally to grant the degrees that the trustees have formally approved. On the other hand, there are also speech acts that fall between Strawson's K-I and K-II extremes. The chair says 'the meeting is adjourned'. His intention is to cause the members of the meeting to stop introducing motions and debating them. If nobody pays any attention, debate goes on, and three more motions are passed, his intentions will surely be frustrated. Equally clearly, his act is intended to 'further a certain [conventional] practice'. It is intended to play the conventional role of adjourning a meeting under sanctions of law or custom.

Strawson is right, of course, that there is a sense in which, for example, the chair's *saying*, at an appropriate time, that the meeting is adjourned 'cannot fail to do so' (1964: 612); that is, cannot fail to adjourn the meeting. No matter what the members go on to do, there is a sense in which 'the meeting has been adjourned' once the chair has spoken, just as after the minister pronounces a pair man and wife they are married, even if they don't act married, and even if everyone else, including those responsible for enforcing the law, fails to treat them as married. 'The meeting is adjourned', said by the right person at the right time, cannot fail to adjourn the meeting, for the simple reason that for the meeting to be adjourned simply *is* for a conventional move to have been made the *conventional* outcome of which would be that no more debate occurs, no more motions are considered and so forth. After the minister or Justice of the Peace has said, at the appropriate time in the appropriate context, 'I now pronounce you man and wife', the couple addressed cannot fail to be married for the unmagical reason that to be married simply *is* for conventional moves to have been made the conventional outcome of which would be that they behaved in certain manners toward one another, were treated by the law in a certain manner, and so forth. But, contra Strawson, the case with K-I acts is exactly parallel. After Mom has said 'Now take out the garbage!' Johnny is under instructions to take out the garbage, whether or not he does so. And after Jane says to John 'Please pass the salt', a request has been made whether or not John complies with it. In these cases too a conventional move has been made. The differences between K-I and K-II illocutionary acts are not quite as sharp as Strawson supposed.

REFERENCES

- Bach, K., and Harnish, R. M. (1979), *Linguistic Communication and Speech Acts* (Cambridge Mass.: MIT Press).
- Grice, H. P. (1967), 'Meaning', *Philosophical Review*, 66/3: 377–88; repr. in Rosenberg and Travis (1971: 436–43).
- Millikan, R. G. (1984), *Language, Thought, and Other Biological Categories* (Cambridge, Mass.: MIT Press).
- (1993), *White Queen Psychology and Other Essays for Alice* (Cambridge, Mass.: MIT Press).
- (1994), 'A Bet with Peacocks', in C. Macdonald and G. Macdonald (eds.), *Philosophy of Psychology: Debates on Psychological Explanation* (Oxford: Oxford University Press).
- (2002), 'Biofunctions: Two Paradigms', in R. Cummins, A. Ariew, and M. Perlman (eds.), *Functions: New Readings in the Philosophy of Psychology and Biology* (Oxford: Oxford University Press), 113–43.
- (2004), *Varieties of Meaning: The Jean Nicod Lectures 2004* (Cambridge, Mass.: MIT Press).
- Recanati, F. (1987), *Meaning and Force: The Pragmatics of Performative Utterances* (Cambridge: Cambridge University Press).
- Rosenberg, J., and Travis, C. (1971) (eds.), *Readings in the Philosophy of Language* (Englewood Cliffs, NJ: Prentice-Hall).
- Schiffer, S. (1972), *Meaning* (Oxford: Clarendon).
- Searle, J. (1969), *Speech Acts: An Essay in the Philosophy of Language* (Cambridge: Cambridge University Press).
- (1989), 'How Performatives Work', *Linguistics and Philosophy*, 12: 535–58.
- Sperber, D., and Wilson, D. (1986), *Relevance: Communication and Cognition* (Cambridge, Mass.: Harvard University Press).
- Strawson, P. F. (1964), 'Intention and Convention in Speech Acts', *Philosophical Review*, 73/4: 439–60; repr. in Rosenberg and Travis (1971: 599–614).
- Warnock, G. J. (1973), 'Some Types of Performative Utterance', in G. J. Warnock (ed.), *Essays on J. L. Austin* (Oxford: Clarendon), 69–89.



## 9 Pushmi-pullyu Representations\*

---

### 1. Introduction

A list of groceries, Professor Anscombe once suggested, might be used as a shopping list, telling what to buy, or it might be used as an inventory list, telling what has been bought (Anscombe 1957). If used as a shopping list, the world is supposed to conform to the representation: if the list does not match what is in the grocery bag, it is what is in the bag that is at fault. But if used as an inventory list, the representation is supposed to conform to the world: if the list does not match what is in the bag, it is the list that is at fault. The first kind of representation, where the world is supposed to conform to the list, can be called directive; it represents or directs what is to be done. The second, where the list is supposed to conform to the world, can be called descriptive; it represents or describes what is the case. I wish to propose that there exist representations that face both these ways at once. With apologies to Dr Dolittle,<sup>1</sup> I call them pushmi-pullyu representations (PPRs) (See Fig. 9.1).

PPRs have both a descriptive and a directive function, yet they are not equivalent to the mere conjunction of a pure descriptive representation and a pure directive one but are more primitive than either. Purely descriptive and purely directive representations are forms requiring a more sophisticated cognitive apparatus to employ them than is

\* This chapter originally appeared in James Tomberlin (ed.), *AI, Connectionism, and Philosophical Psychology*, Philosophical Perspectives, ix (Atascadero, Calif.: Ridgeview, 1995). It is reprinted, with minor revisions, by permission of the author and publisher. The essay was also printed in L. May and M. Friedman (eds.), *Mind and Morals* (Cambridge, Mass.: MIT Press, 1996), 145–61.

<sup>1</sup> And Hugh Lofting, *The Story of Doctor Dolittle; Being the History of His Peculiar Life at Home and Astonishing Adventures in Foreign Parts* (New York: Bantam, Doubleday, Dell, 1922).




Fig. 9.1 Pushmi-pullyu

necessary for these primitives. Purely descriptive representations must be combined with directive representations through a process of practical inference in order to be used by the cognitive systems. Similarly, purely directive representations must be combined with descriptive ones. The employment of PPRs is a much simpler affair.

Perhaps the most obvious PPRs are simple signals to conspecifics used by various animals, such as birdsongs, rabbit thumps, and bee dances. But PPRs also appear in human language and probably in human thought. Illustrations in language are 'No, Johnny, we don't eat peas with our fingers' and 'The meeting is adjourned', as said by the chair of the meeting. Human intentions are probably an example of PPRs in thought, serving at once to direct action and to describe one's future so that one can plan around it. Our inner representations by which we understand the social roles that we play as we play them are probably also PPRs. The natural way that we fall into doing 'what one does', 'what women do', 'what teachers do', and so forth suggests this.

I suspect that these primitive ways of thinking are an essential glue helping to hold human societies together.

I view PPRs from within a general theory of representations developed in other places (Millikan 1984; 1993; 2004). I will start by sketching just a little of that theory (though not nearly enough to defend it).

## 2. *The Background Theory of Representation*

Brentano took the essence of intentionality to be the capacity of the mind to 'intend' the nonexistent. In recent years it has become generally accepted that he was right in this sense: the core of any theory of representation must contain an explanation of how misrepresentation can occur. I have argued that misrepresentation is best understood by embedding the theory of intentionality within a theory of function that allows us to understand more generally what malfunction is. For this, we use a generalization of the biological notion of function as, putting things very crudely, the survival value of a reproduced type of entity. I call functions of this generalized kind 'proper functions' (Millikan 1984; 1993).

Think of the proper function of a type as what it has been doing to account for its continued reproduction. The possibility of misrepresentation is derived from the possibility that a token may fail to perform a function that has been accounting for continued reproduction of its type. In some cases failure to function properly may be statistically even more common than proper functioning. For example, many biological mechanisms fail much of the time to have those occasional propitious effects that have nonetheless accounted for their proliferation in the species. Thus, the eye-blink reflex, exhibited when any object moves too close to the eye, may be triggered uselessly many times for every time it actually prevents foreign matter from entering the eye. In order to proliferate, often a type needs to perform properly only in some critical but small proportion of cases, a proportion that varies widely depending on other factors that enter into the mechanisms of evolution by natural selection.

We can apply the notion of proper function directly to natural languages. Whole sentences are not usually reproduced, but phonemes and words are reproduced, and the syntactic forms in which they are placed are reproduced. They are copied from one generation of speakers to the next and reproduced by the same speaker on various occasions. If some theory of universal grammar is correct, then certain very general grammatical features are also reproduced via the genes. But, more immediately, it is clear that speakers of a given language reproduce words patterned into certain concrete syntactic forms rather than certain other forms because the effects these patterns have had upon hearers, in some critical proportion of cases, are effects the speakers wish to reproduce. These effects are described with reference to the semantic and conventional pragmatic rules of the language.

The proper effects upon hearers of language forms—the proper reactions of hearers to these forms—are also reproduced. This is not necessarily because hearers directly copy one another's reactions, though this may sometimes happen. These reactions are reproduced because hearers, at least often enough, benefit from reacting properly to the sentences they hear. By moving from sentences into belief or into action in conformity with the rules of the language, often enough hearers gain useful knowledge or their actions find a reward; hence they reproduce conformity to these rules. Thus, each generation of hearers learns to accord in its reactions with the expectations of the previous generation of speakers, and each generation of speakers learns to expect those same reactions from a previous generation of hearers—with sufficient frequency, that is, to keep the language forms in circulation. Similarly, the proper reaction of a bird to the song of a conspecific, though not for most species an imitated response, is also reproduced—in this case, genetically reproduced. In both cases the proliferation of representations and the proliferation of proper reactions to them are each dependent on and tailored to the other.

The proper function (or functions) of an expression in a public language may contrast with the function that a speaker intends for it on a given occasion. I have shown how the speaker's intention in use can lend tokens of a language device additional proper functions, as in the case of Gricean implicature (Millikan 1984: ch. 3). These additional

functions are not functions of the public-language forms, however. Indeed, the two layers of function, public and private, can sometimes conflict, as in the case of purposeful lying and certain parasitic uses of language. We will not need to consider this sort of complexity here, but it is well to warn in advance that if there are pushmi-pullyu forms existing in the public language, as I will argue there are, these should not be confused with the well-known phenomenon of Gricean implicature. Public pushmi-pullyu forms are double-aspected on the first layer of function—the public-language layer. It is not merely the user's intention that produces the pushmi-pullyu effect.

To understand what inner representations are (percepts, thoughts), we apply the notion of a proper function not directly to the representations themselves but to the mechanisms whose function it is to produce and to use inner representations. When functioning properly, inner-representation-producing mechanisms produce representations in response to, and appropriate to, situations in which the individual organism finds itself. In humans, these mechanisms are exceedingly complex, including mechanisms of belief- and desire-formation and also mechanisms of concept-formation, inference, decision-making, and action. When the entire system functions properly, the belief-forming mechanisms produce true beliefs and the desire-forming mechanisms produce desires whose fulfillment would benefit the organism. But it is possible also to sharpen the notion 'proper function' in such a way that inner representations themselves are seen to have proper functions, as follows.<sup>2</sup>

Many biological mechanisms are designed to produce alterations in the organism in response to some aspect of the environment, so as to adapt or match the animal to that aspect, hence to serve some further function within that environment. A simple example is the mechanism in the skin of a chameleon that rearranges its pigment so that its color will match that of its environment. This mechanism then serves the further function of concealing the chameleon from predators. Each particular coloring of the chameleon produced is naturally said to have a function too—the same function: concealing the chameleon from predators.

<sup>2</sup> See esp., Millikan (1984: ch. 2; 1993: ch. 11; 1994).

Similarly, any state constituting a stage in a (proper-)functional process, when the shape of the process and hence of the state is determined by input from and as a function of certain features of the environment, can be viewed as itself having proper functions. The proper functions of this state are to help in the production of various further stages in the process to which this state will give rise if the whole system continues to function properly. In this way, specific inner states, even quite unique inner states, unique because they are induced by unique organism–environment relations, may have proper functions. In this way, even such representations as someone’s desire to climb Mount Everest backward, or someone’s belief that persimmons cure mumps, can be considered to have proper functions.

But I have not yet said what any of these proper functions are—what it is that representations, either inner or outer, properly do.

### 3. *Descriptive and Directive Representations*

Elsewhere (1984; 2004) I have defended a proposal to explain how the content of a representation—its satisfaction-condition—is derived. The explanation takes time, and it is different for descriptive and for directive representations. Here I will assume the notion of content, explaining only what I take the difference between descriptives and directives to be.

A representation is directive when it has as a proper function to guide the mechanisms that use it so that they produce its satisfaction-condition. Like a blueprint, it shows what is to be done. Desires are directive representations. To see how this might be so, remember that the proper function of an item can be a function that it is unlikely to perform. Perhaps the sad fact is that an overwhelming majority of our desires never become satisfied. Many (for instance, the desire to square the circle) may even be incapable of becoming satisfied or (for instance, the desire that it rain tomorrow) may be incapable of being satisfied by normal operation of those mechanisms that are designed to help fulfill desires. This has no bearing on the claim that desires are satisfied when things proceed ‘properly’; that is, in the ideal sort of way that accounted

for the survival and proliferation of those integrated systems whose job it is to make and use desires. Surely the job of these systems is, first, to produce desires that would benefit one if fulfilled and then, second, given certain additional propitious inner and outer circumstances, to be moved by them to their fulfillment.

Sentences in the imperative mood are, paradigmatically, directive representations. They proliferate in the language community primarily insofar as they (often enough) help to effect the fulfillment of their satisfaction-conditions.

Unlike directive representations, what makes a representation descriptive is not its function. Rather, the descriptive representation's truth-condition is a condition to which it adapts its interpreters or users in the service of their proper functions (Millikan 1984: ch. 6; 1993: chs. 4–6; 2004: chs. 5–6). It is a condition that must hold if the interpreter's (proper) way of being guided by the representation is to effect fulfillment of the interpreter's functions in accord with design. For example, beliefs are descriptive representations. If my belief that there is an umbrella in the hall closet is to help guide my decision-making and action-guiding apparatus (the belief's 'interpreter') such that it serves its function of helping to fulfill my desires—for example, my desire to keep off the rain—then there needs to be an umbrella in the hall closet. If it is to help me make a correct inference concerning, for example, whether Susan returned my umbrella, one that yields truth not by accident but in accordance with the good design of my cognitive systems, still there needs to be an umbrella in the hall closet, and so forth.

Typical sentences in the declarative syntactic pattern have a descriptive function. Their function is to produce hearer beliefs or, more precisely, to produce true hearer beliefs. For it is only when a certain proportion of what hearers are told is true, and hence is interpreted by them into true beliefs, that they are encouraged to continue to conform to certain rules of interpretation from English (from the public language) into beliefs. One could not learn to understand a language if too large a proportion of the sentences one heard in it were false. Because its function is to produce true beliefs in hearers, the declarative syntactic pattern is descriptive. Roughly speaking, sentences in this pattern affect their interpreters in a proper way only under the conditions that are their truth-conditions.

Earlier I remarked that I was not pausing to explain how the satisfaction-conditions of representations are determined, but it will be important to grasp this much. On the analysis I have given, the satisfaction-condition of a representation is derived relative to its function. The content of the representation turns out to be an abstraction from a fuller affair intrinsically involving an embedding mood or propositional attitude. Put simply, there is no such thing as content without mood or attitude; content is an aspect of attitude. A corollary, as we will soon see, is that it is possible for the very same representation to carry at once two different contents, one relative to each of two different attitudes or moods that simultaneously embed it.

#### 4. *Pushmi-Pullyu Representations*

Consider first a very primitive representation: the food call of a hen to its brood. A proper function of this call is to cause the chicks to come to the place where the food is and so to nourish them. Assume, what is reasonable, that this is the only proper effect that the call has on chicks, the only effect the call has been selected for. Then the call is directive, saying something like 'Come here now and eat!' But it is also a condition for proper performance of the call that there be food when the hen calls, so the call is also descriptive, saying something like 'Here's food now.' (Note that the descriptive and the directive contents of this representation are different.) Assume further, what is again reasonable, that the effect of the call on the chicks is not filtered through an all-purpose cognitive mechanism that operates by first forming a purely descriptive representation (a belief that there is food over there), then retrieving a relevant directive one (the desire to eat), then performing a practical inference, and finally, acting on the conclusion. Rather, the call connects directly with action. *Its function is to mediate the production of a certain kind of behavior such that it varies as a direct function of a certain variation in the environment, thus directly translating the shape of the environment into the shape of a certain kind of conforming action:* where the hen finds food, there the chick will go. The call is a PP representation.



Other examples of primitive PPRs (probably) are other bird calls, danger signals used by the various species, the various predator calls used by chickens and vervet monkeys, and the dance of the honey bee. For example, the bee dance tells at once where the nectar is and where to go. Functioning properly, it produces variation in behavior as a direct function of variation in the environment. Actually, there is evidence that the bee has a map in its head of its environment and that the dance induces it, first, to mark the nectar location on this map (Gallistel 1990).<sup>3</sup> Still, assuming that the only use the bee ever makes of a mark for nectar on its inner map is flying to the marked position to collect the nectar, then the nectar on the bee's inner map is itself a PPR. And it seems reasonable to count a representation whose only immediate proper function is to produce an inner PPR as itself a PPR.

James J. Gibson did not advocate speculating about inner representations. Yet his notion that in perception we perceive certain affordances (opportunities for action) suggests that perceptual representations are PPRs. Think of perceptual representations simply as states of the organism that vary directly according to certain variations in the distal environment. The perceived layout of one's distal environment is, first, a representation of how things out there are arranged—a descriptive representation. It is also a representation of possible ways of moving within that environment: ways of passing through, ways of climbing up, paths to walk on, graspable things, angles from which to grasp them, and so forth. Variations in the layout correspond to variations in possible projects and in the paths of motion needed to achieve them. The representation of a possibility for action is a directive representation. This is because it actually serves a proper function only if and when it is acted upon. There is no reason to represent what can be done unless this sometimes affects its being done. Compare desires, which serve a function only insofar as they occasionally help to produce their own fulfillment. In the case of perceived affordances, action toward their fulfillment is, of course, directly

<sup>3</sup> If the dance says there is nectar at a location where the bee's experience has previously shown there is a large body of water, the bee will be unmoved by the dance (Gallistel 1990).

guided by the percept, variations in the environment, hence in the percept, translating directly into variations in the perceiver's movement.

There are cells in the inferior premotor cortex of monkeys (informally, 'monkey see-monkey do cells') that fire differentially according to the immediate ends (such as grasping small pieces of food with the fingers) of the manual manipulations the monkeys are about to execute and that also fire when the monkeys see other monkeys perform these same manipulations for the same ends (Rizzolatti et al. 1988). Imitative behaviors in children show up extremely early. One infant was observed in the laboratory to imitate facial expressions (opening the mouth, sticking out the tongue, etc.) at the age of forty-two minutes (Meltzoff and Moore 1983). We might speculate, on analogy with the monkeys, that these primitive mechanisms of imitation in children employ PPRs, which picture what the other is doing at the same time that they serve to direct what the self is to do. Compatibly, Jeannerod (1994) cites evidence that imagining oneself performing certain movements and actually performing them involve, in part, the same dedicated area of the brain, hence that picturing what one might do and intending to do it may be two sides of numerically the same representing coin. Indeed, one of his suggestions is that imagining an action without at the same time performing it is accomplished by inhibition of normal connections to motor pathways.

It is important to see that PPRs are not merely conjunctions of a descriptive plus a directive representation. They are more primitive than either purely directive or purely descriptive representations. Representations that tell only what the case is have no ultimate utility unless they combine with representations of goals, and, of course, representations that tell what to do have no utility unless they can combine with representations of facts. It follows that a capacity to make mediate inferences, at least practical mediate inferences, must already be in place if an animal is to use purely descriptive or purely directive representations. The ability to store away information for which one has no immediate use (pure description), and to represent goals one does not yet know how to act on (pure direction), is surely more advanced than the ability to use simple kinds of PP representations.

## 5. *PP Representations in Human Thought*

Organisms often evolve in complexity by modifying less differentiated multipurpose structures into more differentiated dedicated ones. Thus, we would expect beliefs (dedicated to facts) and desires (dedicated to goals) to be a later evolutionary achievement than inner PPRs. On the other hand, if there are purposes that could be served as well or better by PPRs than by more differentiated representations, our first hypothesis should be that that is how these purposes still are served. I think that there are some such purposes.

One obvious hypothesis is that human intentions are PPRs. If intentions are inner representations, surely they are at least directive ones. They perform their proper functions when they issue in the intended actions. But it is also a common and plausible assumption that a person cannot sincerely intend to do a thing without believing she will do it. If one starts in a rather traditional way, assuming that there are only two basic sorts of cognitive representations—purely descriptive ones (beliefs) and purely directive ones (desires)—then whether ‘intending’ must involve harboring a descriptive as well as a directive thought may appear to be a matter of ‘analysis of the concept “intention”’. But if intentions are PPRs, then the dual nature of intentions is no conceptual truth but a biological or neurological truth. And there is reason to suppose they might be PPRs.

Suppose that my brain already harbors, for purposes of guiding my action, a representation of what I am definitely going to do. And suppose there is need to take this settled future into account when making further decisions about what else I can compatibly do. It would surely waste space and introduce unnecessary mechanisms for evolution to duplicate the representation I already have. Better just to use it over again as a descriptive representation as well. Notice, however, that this kind of PPR differs from the kinds I have previously discussed in this way. Rather than functioning as do, say, perceptual PPRs, which map variations in the organism’s world directly into (possible) actions, it maps variations in goals directly on to the represented future world. It differs also in that the contents of the directive and descriptive aspects of the representation are not different but coincide.

A second kind of PPRs that may be fundamental in human thought are primitive representations of social norms and roles. I suggest not that this is the only way humans can cognize these norms and roles but that it may be the primary functional way, and that this way of thinking may serve as an original and primary social adhesive. There are good reasons for thinking that humans and other social animals have designed into them mechanisms leading to the coordination of behaviors.<sup>4</sup> Coordinated behaviors are ones that benefit each individual involved in the coordination given that the others also play their assigned roles. Some of the principles governing evolution of such behaviors are now well known. Allan Gibbard suggests that 'systems of normative control in human beings . . . are adapted to achieve interpersonal coordination' (1990: 64). His ultimate aim is to cast light on the origins and function of the language of ethics and the thought it communicates, including especially the function of normative discussion in originating what he calls 'normative governance'. My project here is less ambitious. My speculations concern only a mechanism for stabilizing and spreading coordinative behaviors that are already in place and for coordinating expectations.

I have in mind two basic sorts of social coordinations. The first might be called *common* norms. They apply equally to all members of a given society: we drive on the right; we speak at meetings only when duly recognized; we wait in orderly queues; we are quiet at concerts; we honor our contracts; we see to our families first, then to our relatives, then to our friends; and so forth. The second might be called *role* norms. These apply to a person only insofar as he or she is filling a certain role: children obey adults while adults direct children; the chair of the meeting calls it to order and so forth, but does not introduce or speak to motions or vote, while the members do introduce motions, speak to them, and vote; pupils raise hands to be called on while teachers speak freely; guests and hosts behave in assigned ways; and so forth.

The examples of norms just mentioned undoubtedly all have coordinating functions. On the other hand, the distinction between common norms and role norms could be applied as well to norms lacking

<sup>4</sup> A good informal summary of these reasons is given in Gibbard (1990).

coordinating functions. For example, not eating peas with one's fingers and not picking one's nose in public may be noncoordinating common norms, whereas wearing a skirt if you are female and trousers if you are male may always have been a noncoordinating role norm. A mechanism whose biological function is to transmit coordinating norms might well have as a mostly benign side effect the transmission of a good number of noncoordinating norms as well. Thus, humans tend to be creatures of convention, exhibiting many patterns of both solo and interactive behavior that are handed on quite blindly, seeming to serve no purpose at all. It may be that our propensity to play games is another side effect of a mentality built to effect coordinations.

The mechanism for stabilizing coordinative patterns of behavior that I propose is simple. It is the capacity and disposition to understand social norms in a way that is undifferentiated between descriptive and directive. What one does (or what *das Mann* does—remember Heidegger?), what a woman does, what a teacher does, how one behaves when one is married or when one is chair of the meeting, these are grasped by means of thoughts, PPRs, that simultaneously describe and prescribe. In the primitive mind these PPRs describe and prescribe what is understood to be The Moral Order: an order taken as totally objective, noninstrumental (absolute), and real, but understood at the same time as stringently prescribed. (In primitive thought, self and others have Sartrean essences with a vengeance.) But it may also be that without the general disposition to think in this way during much of the unreflective parts of our lives, the social fabric would be weakened beyond repair. Yes, I am seriously proposing this as a possible neural mechanism, although supplemented, of course, with more sophisticated mechanisms by which we moderns may also dissect the relevant norms to reveal two faces.

## 6. *PP Representations in Human Language*

If human thought contains PPRs, arguably human languages should contain them as well. Just as they contain forms, on the one hand, whose function is to implant ideas about what is the case and, on the other, to

implant ideas about what to do, one would expect them to contain forms whose function is to implant mental PPRs in hearers. For example, if the inner vehicles of our intentions and our unreflective graspings of social norms are PPRs, it is reasonable that there should be linguistic forms to correspond. There do not seem, however, to be any dedicated syntactic arrangements to do these jobs. And indeed, granted PPRs really exist in thought, this lack of a form of expression dedicated entirely to PPRs may do much to explain their near invisibility. PPRs, I believe, are imparted by use of the declarative syntactic pattern. At least, this is true for English. That is, the declarative pattern has more than one function. Sometimes it is descriptive, and sometimes it expresses a PP mood.

In the PP mood we say, for example, to children 'We don't eat peas with our fingers' and 'Married people only make love with each other'. The job of this mood is to describe and to prescribe, producing at the same time both true expectations and coincident behaviors, the one as a direct function of the other. Notice that the mechanism here is not that of Gricean implicature. Both functions are explicit or literal; the mood is proliferated precisely because it serves both functions at once; both functions are fully conventional.

Strict orders are standardly delivered in the English declarative pattern, which then functions directive: 'You will report to the CO at 6 a.m. sharp'; 'You will not leave the house today, Johnny, until your room is clean'. This use, I suggest, is more than just directive; it is not just another imperative form.<sup>5</sup> Its function is to impart an *intention* to a hearer and to impart it directly, without mediation through any decision-making process; for example, without involving first a desire and a practical inference. This is the PP mood, undifferentiated between directive and descriptive, serving to impart PPRs.

The PP mood may also be used to impart intentions to a group, thus serving a coordinating function. 'The meeting will now come to order', when it functions properly, imparts to each member of the group both intentions concerning their own behavior and expectations concerning the behavior of others. The PP mood statement in the university catalog, 'Professors will hold office hours every day during registration

<sup>5</sup> This is a correction of the remarks on army orders in Millikan (1984: ch. 3).

week' informs both students and faculty and imparts intentions to faculty.

I would like finally to introduce two more kinds of PPRs that I believe occur in public language but that will require somewhat more lengthy discussions. The first is a variety of explicitly performative sentence. The second are terms expressing 'thick concepts'.

## 7. *Performatives*

Many performative utterances, I believe, are in the PP mood.<sup>6</sup> When viewed incorrectly as simple present-tense descriptives, they are puzzling in that they seem to create facts *ex nihilo*, making something to be the case simply by saying that it is the case. Suppose that the chair of the meeting says 'The meeting is adjourned', but nobody pays any attention and three more motions are debated. There does seem to be a sense in which these motions will have been debated not just after the meeting was *called* adjourned but after it *was* adjourned; that is, a sense in which 'the meeting is adjourned' is self-guaranteeing. Similarly, once the chair has said 'The meeting will now come to order', the meeting has been brought to order even if everyone keeps talking loudly. And if the right person says to the right couple at the right time 'I now pronounce you man and wife', then these two are man and wife, even if they do not act that way and even if nobody, including the legal authorities, is disposed to treat them that way. On the other hand, if nobody pays any attention when the chair says 'The meeting is adjourned' or 'The meeting will now come to order', there is also a tug that says these sentences somehow were not true. Let me try to explain why this tension occurs.

Many conventionally or legally molded patterns of activity allow or require the making of, as I will call them, conventional or legal 'moves',

<sup>6</sup> The performative sentences most obviously referred to in the following discussion are those that perform what Strawson (1964) termed 'essentially conventional illocutionary acts' or, more accurately I believe, 'K-II illocutionary acts' as these are described in Ch. 8. The relation of these performatives to those that perform explicit 'K-I illocutionary acts' is interesting but too involved to pursue in a general discussion of this sort.

under designated circumstances, by participants' playing designated roles—for example, performing a marriage ceremony, making a move in a game, appointing someone to a position, making a motion in accordance with parliamentary procedure, or legally closing a road. These are acts that, once performed, place restrictions on what may or must follow after if things are to accord with convention or law. Moves of this sort may or may not be made with the aid of articulate language forms. For example, in some contexts a bid may be made by saying 'I bid', but in others merely by raising a finger, a vote may be made by raising a hand, and so forth. Most such moves are themselves in essence PPRs.<sup>7</sup> Their proper function is to channel behaviors that follow so that they take certain forms and not others (directive) and to coordinate expectations accordingly (descriptive). Functioning properly in the usual way, they produce inner PPRs in the various participants, which guide them in coordinating their expectations and actions.

The implications of such moves for ensuing behavior are often quite complex, involving complicated mandates or the limiting of options for a variety of affected people playing a variety of roles. For this reason, these implications often could not easily be spelled out in an explicit formula, for example, by saying something simple like 'The meeting will now come to order'. Imagine the minister's trying to fill in all the blanks in this formula: 'You, Jane, and you, John, will now . . . and the guardians of law will now . . . and your parents and friends will now . . .', and so forth, spelling out all that marriage entails. However, there usually are names for conventional or legal moves, or for the situations constituted by their having been performed: 'bidding six diamonds', 'checking the queen', 'being married', 'being chairman', 'appointing a chairman', 'voting for a candidate', 'a road being legally closed', 'making a motion', and so forth. These names are not names for the 'shapes' of the moves—they are not names like 'raising one's finger' or 'raising one's hand'. Rather, they are names for descriptions that classify moves according to their conventional outcomes—by what sorts of things follow after them in accordance with convention or law.

<sup>7</sup> I say 'most' because some conventional patterns have no functions at all (see Ch. 1; Ch. 8).



Now, there is an obvious and very general meta-convention by which one may use the name of any move in order to make that move, for assuming that the conventional 'shape' that the move has is merely conventional, as it generally is, any other shape might be substituted for this shape so long as everyone understands what move is being made. Saying that I perform the move will generally be enough to perform it, then, granted that the vehicle that usually performs it is arbitrary, and granted that I am the fellow whose move it is to make. For example, even 'I move my queen to Q5' may be quite enough to perform that move in chess, especially if circumstances are such as to make an actual board move awkward or impossible (such as chess by mail or while also trying to get dinner).

In English, present-tense declaratives are used in making such moves: 'I bid three diamonds', 'I pronounce you man and wife', 'I move that the meeting be adjourned', 'The meeting is adjourned', 'This road is legally closed', and so forth.<sup>8</sup> This use of the declarative pattern proliferates because it serves the function of producing the conventional outcomes of the moves named, channeling ensuing activities (a directive function) and coordinating relevant expectations (a descriptive function) at the same time. The proper function of the declarative in each of these cases is thus exactly the same as the function of the move named or described; indeed, the move *is* the utterance, in the right circumstances, of the sentence naming it, just as raising one's hand in the right circumstances is voting. It is a matter of convention that one can vote by raising one's hand, and it is a matter of (very sensible) convention that one can make almost any conventional move by embedding its name in this way in the declarative pattern. These embeddings are called 'performative formulas'. Such performative formulas are PPRs. They produce conventional outcomes and coordinate expectations accordingly, standardly by producing relevant inner PPRs in participants.

<sup>8</sup> Some have argued that performatives function by conversational implicature. However, explicit performatives like 'I order you to', 'I promise that', and so forth are lacking in some languages, 'particularly those with a more developed system of sentence types or those spoken in societies that seem to have less cultural need for formulaic discourse of the kind represented by performative sentences' (Sadock 1988: 186).

But there is another detail. One can, of course, make a conventional move with a performative formula that names that move only granted that one is the right person in the right circumstances to make that move. I cannot, for example, adjourn a meeting by saying 'The meeting is adjourned' unless I am chair and we are in the right kind of a meeting and the right time has come. Another function of the performative formula—another thing it does when functioning properly—is to produce true beliefs that the named conventional move is in fact being performed. In this respect, the performative is just (is also) an ordinary descriptive. It is true if it is in fact being uttered by an appropriate person in appropriate circumstances; that is, if the move it claims is performed is in fact being performed. Nor does this hang on whether it in fact effects its conventional outcome.

Suppose that I put up a sign on my road saying 'This road is legally closed'. If I have reproduced this token of the declarative pattern on the model merely of descriptive sentences I have previously heard rather than copying my use from the cultural pattern of performative uses, then it is not, of course, a sentence in the PP mood at all. It is merely descriptive, and, assuming the town has not in fact closed my road, it is plain false. If I have reproduced it instead on the model of PP sentences I have previously heard, perhaps supposing myself to have the legal right to close my own road in this way, then it is in the PP mood. Still, just as if it were a simple descriptive, it is, minimally, false. It is not true that my road is legally closed.

Integrating this now with the previous theme, besides this last sort of truth condition a performative sentence of this kind also has a directive satisfaction-condition and a second truth condition as well. It directs that a certain conventional outcome should ensue, and it induces expectations to accord with this. Returning to 'The meeting is adjourned', one of its truth conditions is fulfilled by the fact that it is the chair who says it and at an appropriate time. Whether or not three more motions are debated afterward, once the chair has said these words, the meeting is in fact adjourned. There is a second truth condition, however: the proceedings should then actually draw to a close. Only if this happens will the coordinated expectations that it is the proper job of the performative sentence to induce be true expectations. Last, like any PP mood sentence

whose job is to impart intentions, it has the meeting's actually drawing to a close as a directive satisfaction condition as well.

## 8. *Thick Concepts*

I wish now to speculate about one last example of PPRs in public language: sentences that contain words expressing 'thick concepts'. Thick concepts are concepts such as *rude*, *glorious*, or *graceful*, which seem to describe a thing and to prescribe an attitude toward it at the same time. Recently there have been absorbing and complex discussions about these concepts that I cannot enter into here. But I would like to make a suggestion that might be fruitful to explore.

I hesitate to take any strong position on thick concepts in part because I am not sure, within the framework I have been using, how to understand the nature of those states or dispositions that ordinary speech calls 'attitudes'. (Philosophers' 'propositional attitudes' of belief and desire are not 'attitudes' in ordinary speech at all.) It is not likely, for example, that an adequate description of the attitude involved in thinking something rude could be given by reference only to the notions 'directive' and 'descriptive'. But let us suppose this much: if the function of a term in the language is in part to induce an attitude, then this part of its function is not entirely descriptive. That is, attitudes are other than or more than mere (descriptive) beliefs. If I can make plausible, then, that one function of a term expressing a thick concept is to produce an attitude but that these terms are also descriptive, I shall have shown at least that these terms are close kin to PP representations.

Suppose that certain (perhaps highly disjunctive) configurations of primary qualities tend to produce in us certain attitudes toward their bearers when perceived or contemplated. This might be due to native dispositions or the influence of culture. Secondary qualities are traditionally thought of as powers to produce sensations, but powers to produce attitudes are perhaps similar enough to be called secondary qualities too—'attitudinal secondary qualities'. Secondary qualities are such only relative to a kind of perceiver or, in this case, a kind of reactor.

Nonetheless, relative to a certain group or class of reactors, attitudinal secondary qualities are objective properties.

My suggestion is that words expressing thick concepts may, first, describe attitudinal secondary qualities relative either to our species as a whole or to the culture of speaker and hearer. That is, declarative sentences using these words attributively will not serve their proper functions in a normal way unless the objects of attribution do indeed have certain attitudinal secondary qualities relative to a community encompassing both speaker and hearer. But, second, this is because their proper function is a directive one: to produce in the hearer the relevant attitudes. Their function is to cause the hearer to take these attitudes toward these items. That is, these words continue to produce in hearers these attitudes toward designated objects only because the attitudes induced turn out, in a large enough proportion of cases, to be independently 'true' by the hearers' own lights. The objects described are such as actually to produce, on direct inspection or given a more detailed description, those attitudes. Words expressing thick concepts thus face two ways at once, describing their subjects and at the same time inducing attitudes toward these subjects. Indeed, perhaps the inner representations induced by these words themselves face two ways at once. Perhaps they are inner PPRs.

## REFERENCES

- Anscombe, G. E. M. (1957), *Intention* (Ithaca, NY: Cornell University Press).
- Gallistel, C. R. (1990), *The Organization of Learning* (Cambridge, Mass.: MIT Press).
- Gibbard, A. (1990), *Wise Choices, Apt Feelings* (Cambridge, Mass.: Harvard University Press).
- Jeannerod, M. (1994), 'The Representing Grain: Neural Correlates of Motor Intention and Imagery', *Behavioral and Brain Sciences*, 17: 187-245.
- Meltzoff, N., and Moore, K. (1983), 'Newborn Infants Imitate Adult Facial Gestures', *Child Development*, 54: 702-9.
- Millikan, R. G. (1984), *Language, Thought and Other Biological Categories* (Cambridge, Mass.: MIT Press).

- Millikan, R. G. (1993), *White Queen Psychology and Other Essays for Alice* (Cambridge, Mass.: MIT Press).
- (1994), 'A Bet with Peacocks', in C. Macdonald and G. Macdonald, *Philosophy of Psychology: Debates on Psychological Explanation* (Oxford: Oxford University Press).
- (2004), *Varieties of Meaning* (Cambridge, Mass.: MIT Press).
- Rizzolatti, G., Carmada, R., Fogassi, L, Gentilucci, M., Luppino, G., and Matelli, M. (1988), 'Functional Organization at Area 6 in the Macaque Monkey ii. Area F5 and the Control of Distal Movements', *Experimental Brain Research*, 71: 491–507.
- Sadock, J. M. (1988), 'Speech Act Distinctions in Grammar', in *Linguistics: The Cambridge Survey*, ii (Cambridge: Cambridge University Press).
- Strawson, P. F. (1964), 'Intention and Convention in Speech Acts', *Philosophical Review* 73/4: 439–60.

## 10 Semantics/Pragmatics: (Purposes and Cross-Purposes)\*

---

### 1. *About Purposes*

Products emerging from histories of natural selection under consistent selection pressures can be described as having ‘natural purposes’: purposes to perform the functions or to produce the effects for which they were selected. A number of different kinds of selection mechanisms generate natural purposes, and selection operates on a number of different kinds of replicators. Besides genes and behaviors for which genes are directly responsible there are behaviors established by conditioning, behaviors learned by consciously intended trial and error, action alternatives selected through trial and error in thought (‘Popperian selection’), and a variety of cultural replicators of the kind Dawkins called ‘memes’ (1976/89: ch. 11)—items that are replicated by people copying one another’s behaviors and products. Purposes derived from natural selection are often contrasted with human purposes. I have argued that this is a mistake, that these purposes are not of a fundamentally different kind (1984; 2000: ch. 1; Ch. 5 above). Even those human purposes that are determined by explicit goals and conscious intentions are merely sophisticated forms of natural purposes.

\* Much of the first half of this chapter has been revised and adapted, with the kind permission of the Hegeler Institute, from ‘Purposes and Cross-purposes: On the Evolution of Language and Languages’, *Monist* 84/3 (July 2001), 392–416. I am grateful to Robin Carston for a helpful reading of an earlier draft of this ‘second edition’. I use ‘Purposes and Cross-purposes’ as a subtitle here because in a footnote to chapter 10 in *Varieties of Meaning* (2004) I promised, I now think inadvisedly, to present the material on understanding implicatures and nonliteral usages that appears here (Sect. 4) in a chapter named that.

Representations of goals, explicit intentions, and so forth, have fulfillment of their represented ends as 'derived proper functions' (Millikan 1984: ch. 2; 2002), functions that are derived from the complex biological functions of the cognitive and conative mechanisms that produced them, which mechanisms were designed by natural selection to vary their productions to accord with relevant variations in inner and outer conditions and circumstances.

I will be concerned here with interactions between two broad kinds of purposes. First are the purposes of human behaviors. These form a variety, ranging from directly genetically determined purposes such as the purpose of the eye-blink reflex, through purposes that result, for example, from subtle or less subtle conditioning or that govern the fine-tuned intermediate portions of automatized skills, to purposes defined by explicit intentions and represented goals. Second are the purposes of conventional public-language forms such as words, idioms, syntactic forms, and so forth. Public-language forms are 'conventional' in the sense described in the first chapter of this book. They are copied or 'reproduced' in a language community owing to precedent determined by historical accident, rather than owing to properties that make them more intrinsically serviceable than other forms would have been. That is, they are relatively arbitrary. Conventional linguistic forms are memes of a particular sort, memes that are selected for because they are serving *coordinating* functions (see Chs. 1–5). Their proliferation or 'survival' depends on a sufficiency of occasions on which they promote *both* a speaker's *and* a hearer's interests. Their coordinating functions, their memetic purposes, determine their conventional meanings of both the first and the second kinds discussed in Chapter 3—'stabilizing functions' and 'semantic mappings'. Given this, my topic, broadly conceived, is the relation of semantics to pragmatics and, with post-Gricean pragmatics in mind, the question of what kinds of speaker purposes and hearer understandings are involved during normal semantic, and during normal pragmatic, linguistic processing.

I'll start Section 2 with a brief review of the way in which conventional linguistic forms acquire memetic functions and a preliminary discussion of the distinction and relation between public-language

semantics and speaker meanings. Then, in Section 3 I will question the sharpness of the distinction between these two kinds of meaning and introduce a role that context may play in thoroughly conventional forms of communication. Last, in Section 4 I will explore the ways in which mature hearers and also children arrive at understandings both of public-language meaning and of speaker meanings, keeping the contemporary question in the foreground whether hearers have to possess some sort of 'theory of mind' in order to do this.

## 2. *Linguistic Purposes*

For a mechanism of selection to contribute to the emergence of items with natural purposes, the method of replication it depends on has to be quite accurate. Copies of copies of copies must continue to be like the originals in relevant respects with very few exceptions. Dawkins (1976/89) calls this 'fidelity'. In the case of human language, at least two kinds of dedicated filters, constraining variety and inhibiting drift in replication, seem to exist. First, there is evidence that the human auditory systems are specifically designed to accomplish efficient mastery of the prosodic and 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 and what as a different pattern. They define the basic same-different scheme for a spoken language, what counts as correct reproduction of an element such as a word or a sentence, thus helping to ensure fidelity in linguistic replication. And they enable a learner to know in advance what aspects of the speech signals produced by herself or others will be the instrumental aspects, aspects that matter to sameness and difference of meaning. Second, universal grammar may be a mechanism helping to effect faithful reproduction of linguistic forms. To posit universal grammar is to posit 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, but it does not, of course, determine what the replicated linguistic forms shall mean, what their functions shall be. Their functions are determined by whatever accounts for their continued proliferation and this, with few exceptions, is some service that they are performing for cooperative speaker–hearer pairs. Correlatively, the replicators that form the basis of a language in use are not utterances only. They 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. This theme is spelled out at some length in my (1984: chs. 3–4, 10–14), so I will just give a couple of easy examples here.

Consider, for any language, the syntactic form that gets labeled ‘indicative’. 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 of its central functions is this: it effects production of a *true belief* having whatever propositional content the various other aspects of the sentence are designed to impart.<sup>1</sup> 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 she heard so as sometimes to extract genuine information from them would soon cease to form beliefs on their basis. She might first try out other interpretations of the form, and different interpretations of other linguistic elements used with it, but eventually she 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, a purpose, of the indicative form itself, whether or not a particular speaker and/or hearer have as their own purpose to use

<sup>1</sup> This function is, of course, highly relational. See my (1984: chs. 2, 4, app. A).

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—advice, instructions, directions, friendly requests, sanctioned directives, and so forth—imperative syntactic forms would either become obsolete or change their functions. Giving a different sort of example, to help to produce representations of elephants in hearers' minds is a purpose of the word 'elephant', further aspects of its memetic function depending on linguistic context.

Besides their conventional or 'stabilizing' functions, tokens of language forms in use also have functions/purposes that are derived from the purposes of the individual speakers who use them. A tool designed by a person for a purpose has that purpose as its own purpose. It has a 'derived' function, derived from the purpose of the person who made it (1984: ch. 2). Similarly, a token of a linguistic form that has been produced by a speaker for a purpose has a derived function, derived from the purpose of the speaker in speaking. Thus, it has two kinds of purposes, a memetic purpose and the speaker's purpose, and these may coincide, or they may differ or even conflict.<sup>2</sup> Indeed, following tradition, it seems clear that failure of these two kinds of purpose to coincide is a common occurrence. In the case of metaphorical use, sarcastic use, or the use of other figures of speech, there is actually a conflict between stabilizing function and speaker purpose. The 'literal meaning' is not what the speaker means. And in uses of the kind that Grice labeled 'conversational implicatures' what the speaker means either conflicts with the stabilizing function of the form or has some additional purpose beyond.

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) for many of the coordination problems of communication. No single participant can gain by unilaterally changing the basic strategies or rules by which she communicates. A settled language community is not easily invaded by incompatible policies of

<sup>2</sup> It is very common for a behavior or product of behavior to have more than one source of function or purpose, and not uncommon for these purposes to conflict. Millikan (2004: ch. 1) contains an extensive discussion of this point.

language use. Still, because the forms of a specific language are arbitrary within very broad limits, they remain subject to slow copying drift. Slow changes do occur both in meanings and in phonological and syntactic structures. More germane to my intended topic, 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 may later become conventional.

### 3. *The Blurry Line between Semantics and Pragmatics*

If stabilizing function and speaker purpose for a linguistic form conflict often enough the result will be either a change in function or the addition of a new function. If the speaker's purpose is cooperative, the hearer understands the purpose, and if this speaker–hearer pattern is reproduced often enough the new use will become conventional. It has become conventional when it starts being copied from one speaker–hearer pair to another rather than reinvented each time it appears, so that it wouldn't occur, or occur nearly so often, were it not for this precedent. Conventionality is not a matter of linguistic rules having suddenly materialized in some obscure quarter or medium. Nowhere are there any rules written or any 'codes' laid down. Conventionality is merely a matter of the likelihood of persons saying and understanding linguistic forms in certain ways had they not experienced their being used in those ways before (see Ch. 1). Conventionality, then, is a matter of degree, and if the distinction between semantics and pragmatics is interpreted as the distinction between what is conventional in language use and understanding and what is not,<sup>3</sup> it follows that the semantics/pragmatics distinction is also a matter of degree. I propose now to expand this theme into several dimensions. Implications for the theory of language understanding will be discussed in Section 4.

Consider, first, the process by which a fresh metaphor may become trite and eventually quite dead. For example, the metaphor 'dead

<sup>3</sup> I recommended this position in my (2004: ch. 11).

metaphor' is fairly dead; were it not, 'fairly dead' would come through as an oxymoron. In the beginning, some people may copy the use of an easy metaphor from friends, while others arrive at it spontaneously. Then there may be a long period, perhaps even encompassing the whole rest of its history, when the form is often used *both* because its new meaning is understood and copied *and* because its old meaning continues to be understood. Examples of this usage will belong to two lineages—will exemplify two different least types—at once (see Ch. 3). The form is copied with both the old and the new uses in mind. Both are resonating, as it were (or perhaps as it is), in the connectionist net. That is what gives many usages their strong color long after they have ceased to be merely metaphorical. Establishment of a new meaning as an independent literal meaning might be defined as occurring when the new meaning would continue to proliferate even if the old meaning were to die out. But although this may be a theoretically clarifying idea, it is hard to see how one might, in practice, gain evidence for a usage having this status, unless, of course, the old usage actually does die out (as often happens).

Besides metaphors and other figures of speech that may die and become literal quite gradually, there are other factors that have a broadly blurring effect on the semantics/pragmatics distinction. I have mentioned that certain parameters of the reproduction of linguistic memetic patterns are constrained by the conventions of phonological structure and syntax, perhaps even aided by innate tendencies. But certain other important aspects of linguistic reproduction seem to be left unfettered, allowing different speakers and hearers to interpret and to copy what they hear somewhat differently. At least these three aspects of speech reproduction seem to allow some slippage:

- the length of the speech-stream segment that is copied;
- the exact effects owing to which a segment is copied, thus what its function or meaning is understood to be;
- the degree of embedding in extralinguistic context that may be copied along with the segment.

I will argue that none of these parameters has a determinate setting for the proliferation of language forms. Thus, there may be variation in

how different persons process the same forms; indeed, in how the same person processes them on different occasions. This adds to the reason why a clear distinction often cannot be drawn between an expression's having one sense or several senses, between its being used in more than one literal sense, or in one literal sense plus various extended or figurative senses, between what has been said and what merely conveyed. The distinction between semantic and pragmatic phenomena rests merely on statistical regularities in psychological processing, and these regularities often are not well defined. Moreover, when different patterns of processing achieve very much the same practical results, there is no pressure for uniformity in processing.

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 even 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, probably these last would 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 what constitutes speaking 'idiomatically'. Beginning at 18 months, children learn perhaps ten words a day, steadily, until well into their teens (Bloom 2000). There is no reason to suppose their capacity for learning conventional phrasings is any less 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 been 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'. A result is that 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 'Anchor's Away', and from 'God be with you' to 'Goodbye'.

Conventions of phrasing are sometimes called 'conventions of language use' as opposed to 'conventions of meaning',<sup>4</sup> 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. Thus, given the prior conventions governing the individual components of the phrase 'This in N', there is nothing arbitrary about what this phrase means when used in answering the phone. Meanings of such phrases are readily derivable from living principles governing their parts. Still, they have something in common with half-dead metaphors, for their meanings are derived alternatively or, often, simultaneously in two different ways, (1) compositionally and (2) from holistic reproduction. It can happen then that these two sources of meaning eventually come apart, the holistic meaning floating free of the compositional structure and acquiring an independent phonology, as in 'Goodbye' and 'Bye'.

Where a chunk's meaning has not pulled completely 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, even for the same person,

<sup>4</sup> Compare (Searle 1975) and (Morgan 1978).

from occasion to occasion. 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, its lineage will remain unbroken through these differences. It will have only one sense in the public language.

Similarly, when words or longer expressions are frequently used in extended ways, ways with a new and different meaning, as in metaphors, euphemisms, or other figures of speech, different people, whether as speakers or hearers, may also hear the expression differently. Indeed, the same person may hear it differently on different occasions. Sometimes it carries strong echoes from its original literal usage, sometimes less strong, and sometimes none. Tokens that are produced with echoes in mind are hybrids, produced by the crossing of two gradually separating lineages. They are often reproduced by 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. Expressions of this intermediate kind have two conventional 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.

The same analysis applies to half-dead conversational implicatures. Tokens of the much celebrated 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 literal requests. But they are also literal questions, which is why they have a tentative rather than a demanding feel.

The lesson seems to be that we need not always choose between this being the literal meaning and that being the literal meaning. Ockham's razor employed to prohibit proliferation of semantic meanings can be as useless as it is for prohibiting the proliferation of living species. Using a closer analogy, it can be as useless as trying to prohibit the memorization of those addition and multiplication 'facts' that we use from memory rather than calculating again each time they are needed.

On the other hand, complete separations may eventually occur between lineages that had earlier reproduced tokens jointly. A classic case is 'The dog went to the bathroom on the living-room rug' in which 'went to the bathroom' floats entirely free of its original moorings (Morgan 1978).

Now consider the question, for the sake of what effect is a linguistic segment replicated? That is the same as asking what its linguistic function, its meaning, is. What is to prevent different people from noticing somewhat different effects, or from generalizing in somewhat different ways to new cases? How broad or narrow, abstract or concrete is a word's meaning? How many distinct senses does it have? None of these questions appears to have answers that are bounded in as sharp a way as are standard pronunciation and grammar.

Consider the colors of things; for example, consider red hair. If a dress were that color, it would never be called red. Our cat that everyone calls orange exactly matches our wooden 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 it true that what is orange for a cat is 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, or third *versus* the fourth dimension? How many different semantic senses does 'clear' have when I clear the table, clear the ground, when the water is clear, the coast is clear, the bell is clear, and the message is clear? Is the term 'lineage' used literally when applied to chains of copied words as in Chapter 3 above, or is this an extended use? Suppose that you have 'gone around' a squirrel in the sense that you have circled from north to east to south to west of a tree it is climbing, but suppose it kept ducking behind so that you never saw its backside.<sup>5</sup> Is this an extension (an abbreviating) of the meaning of 'going around', or does 'going around' have two different meanings one of which covers this case, the other not?

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, and hence where the copying chains have begun to

<sup>5</sup> The example is from William James.



meander in 'new' directions, is anywhere written. Similar remarks may go for the seasoned question whether such words as 'some' and 'two' have more than one semantic meaning (at least some/only some, at least two/exactly two) or whether the shift from one of these meanings to the other is pragmatic. Similarly, for whether 'She had a baby and got married' carries information about temporal order semantically or only pragmatically. There is no reason to suppose that the various lineages concerned here are cleanly formed into separate least types, indeed every reason to suppose they are not. Ockham's razor can't prevent the multiplication of actual complexities.

Last, consider whether what gets copied is expressions or expressions-in-contexts. Where the relevant contexts are linguistic, this question may merge 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 only to some past stretch of time, that stretch being calculated from pragmatic 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 recurs 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, but if Whitehead and Russell wrote *Principia Mathematica* it doesn't follow that Whitehead wrote *Principia Mathematica* (Harnish 1976). Does 'A and B' have two meanings, A and B combined versus A and B separately, or does it have one vague meaning that must be specified further pragmatically? The chunking possibility is that subject-plus-verb-ascribing-responsibility is copied whole so as always to mean the subject has sole responsibility. Then the difference 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 seems 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 somehow *imposed* on all speakers of English, nor even that it is determinate for individual speakers exactly which aspects of usage they themselves copy. The same phrase might be copied on different models at different times, or copied from two models at once. Nor will it matter to the hearer exactly which aspects of language use are the ones that were copied. 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 refers to an indefinite time-span, 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 event is to be interpreted as of a recurring kind 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. Speakers sometimes purposely mislead or make jokes by producing tokens that would have one meaning if derived one way but another if derived another.

Not just position in linguistic context but position in extralinguistic context may also be copied by speakers. Saying 'Hit me!' when playing blackjack to ask the dealer for another card is copied whole. That is, what is copied, hence what has a conventional meaning, is *saying-'Hit me!'-while-playing-blackjack*. Similarly, *saying-'Break-a-leg!'-to-an-actor-before-a-performance* is copied whole by speakers in order, conventionally,

to wish the actor luck. These phrases-in-context are reproduced items, replicators, each with a special memetic purpose or meaning. Contrast the case of 'This is N' used to introduce oneself on the phone. Its use in a particular context is also copied, but saying 'This is N' does not have a separate meaning in that context. It does not mean what it means because it is said on the phone.

Perhaps the simplest example of a conventional use of context to fill out meaning is putting signs or labels on things: 'First Prize' or 'Weight limit 20 tons'. Demonstratives and indexicals are obvious examples of language forms with meanings that depend on context. They also illustrate well how the conventional shades into the pragmatic. The relations that demonstratives bear to their referents in the external environment or in the discourse environment seem often to be conventional but may also be improvised. Use of a pointing gesture along with a demonstrative may typically be copied, different cultures exhibiting different forms of pointing and perhaps different forms for different occasions. There is use of the whole hand, the first finger, the middle finger, the head, and even the lips. Besides pointing are many other ways to establish joint attention as one refers with demonstratives, ways of making salient or using the natural salience of objects in the external or discourse environment. 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 memetic function, slowly entering the semantic dimension, perhaps also becoming stereotyped or 'conventionalized'. 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, so not easily open to a priori inspection or argument. Nor should we expect to find this distinction clearly within individual idiolects, for there is no reason to suppose that the same person always processes the same form the same way nor, given contemporary connectionist views, that these different ways of processing always remain distinct within individual minds.

Like the referents of demonstratives, the domains of quantifiers are also understood relative to context, and again it is an intrinsically blurry question how much the use of context is improvised and how much it follows conventional forms. The domain of a quantifier is often

whatever domain is currently the object of joint attention, either naturally so, or because the speaker has done or said something that conventionally makes it so. With both demonstratives and quantifiers, the speaker herself or certain facts about her, along with 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 stereotyped and conventionalized, cannot be a very determinate matter. Nor can it be determinate exactly how large the copied language-plus-environment chunks are, hence exactly which chunks have semantic meaning. Once again, it cannot be determinate where the semantics ends and the pragmatics begins. When misunderstandings occur, often there is nothing to determine whether the speaker has said or conventionally indicated something wrong, or merely meant something her hearer failed to interpret.

Nor does this fuzziness between semantics and pragmatics raise any special problems for the theory of language interpretation. It has been common to presume a wide divide between the way semantically conveyed information is interpreted or 'processed' by a hearer and the way pragmatically conveyed information is processed. Semantically conveyed information is said to be merely 'decoded' while pragmatically conveyed information must be 'recovered' by inference about the speaker's intentions. In Section 4 below I will explore the role that consideration of speaker intentions plays or does not play in language understanding. But it is important to observe first that semantically conveyed information is *never* simply 'decoded'.

The deep reason for this is that, on the present analysis, there can be no strictly dedicated forms in a language. Language is just a raggedy collection of reproduced speaker-hearer patterns having various origins and independent histories. There is nothing to prevent the same physical sign pattern from emerging from the employment of separate linguistic conventions; indeed, it seems inevitable that linguistic conventions should sometimes accidentally cross. It is true that *political* 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. For example, a state law might require that *Robert's Rules of Order* be followed in certain

political meetings, hence 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 natural languages; there are no linguistic rules. No idiom is dedicated *de jure*. Many may be unique *de facto*, but only within a certain tradition. Other traditions may always develop that crossover and intercede.

Glance at the section on equivocation in any informal logic text to find dozens of amusing examples of crossing conventions. Simple cases involve homonyms and syntactic ambiguities. For a different kind of example, notice that although the whole configuration, '*Hit me!*'-*said-while-playing-blackjack* has its own special memetic meaning, not every instance of someone saying 'Hit me!' while playing blackjack need be a replication of this pattern. Someone might say 'Hit me!' playfully or perhaps defiantly while playing blackjack, having composed this phrase compositionally on the model of the ordinary words and syntax involved. Then, although 'Hit me!' would be uttered while playing blackjack it would not be a member of the least type that means *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 sounds of that kind in that context are from that lineage. It does not force that configuration of sound and context to be everywhere dedicated to that function. Wider context may always suddenly be needed for interpretation—needed in order to make it clear from which lineage of tokens this token has been copied, hence which is its true memetic function. Understanding the semantics of a linguistic form never reduces merely to decoding.

#### 4. *Understanding What Speakers Intend*

I have suggested that there may be much more that is conventional in language use than has usually been supposed. On the other hand, the fact that a use is conventional does not imply that its interpretation is automatic; the fact that the speaker uses a linguistic form in simple

accord with its semantic meaning does not imply that it requires only to be 'decoded'. This is particularly obvious for words and phrases that sound alike but that are composed of semantically separate least types. Consider, for example, the 'two' that means at least two and the 'two' that means exactly two. It may be that these two 'twos' compose separate least types. But the fact that tokens having these different semantic values were derived from separate lineages with separate histories isn't written on their faces. The hearer cannot tell from which lineage one has descended just by its sound. Indeed, figuring out their semantic values seems, ironically, to be a pragmatic problem of exactly the same kind as would exist if the difference between them was merely pragmatic; for example, if one was always derived independently by Gricean implicature from the other. The results we have turned up so far then are merely negative. We have seen reason to avoid arguments about whether this or that understood meaning of a linguistic form is determined directly by its semantic value or instead is an extended meaning, the product of a figure of speech, or of an implicature. These distinctions are intrinsically vague, difficult to define and to measure. And it also seems that not much turns on them. The interesting question for pragmatics, how a hearer manages to interpret the linguistic forms she hears, is not answered merely by the discovery that more forms than we had supposed may be conventional.

According to the Gricean tradition, during normal linguistic communication the hearer, H, recognizes that the speaker, S, has a certain intention in speaking; say, the intention to get H to believe something or, perhaps, to do something. H also recognizes that S intends H to recognize S's intention and, if sufficiently trusting, H proceeds, for this very reason, to fulfill S's intention. Given this analysis, the basic problem posed for H is to figure out what it is that S intends H to do. In my (1984: ch. 3) I argued that this Gricean analysis is very implausible if taken at face value as requiring that speakers and hearers harbor multiply embedded mental representations of one another's mental representations during normal conversation.<sup>6</sup> One reason I gave was that it seemed to have been shown by students of child

<sup>6</sup> It is hard to see how Gricean intentions are to be understood other than as explicit mental representations of this sort, since explicitness would seem to be necessary for

language that children younger than about four, although fairly proficient in the use of language, don't yet have concepts of such things as beliefs, desires, and intentions. Since then there has been much more discussion and considerable experimental work done on the question of when normal children, and also autistic children, Down-syndrome children, deaf children, and blind children, acquire a 'theory of mind', and there seems to be very solid evidence that this depends both on the prior development of skill in handling dependent clauses ('X says, sees, is afraid, . . . that p') and on enough social interaction involving discussion of mental attitudes, rather than acquisition of language depending on already having a 'theory of mind'. Particularly interesting is that deaf children who are otherwise normal often seem not to acquire 'theory of mind' until 15 or 16 years old unless they have been lucky enough to learn sign language from native signers in infancy.<sup>7</sup> But, on the other hand, strong arguments have been made that the fast and accurate way small children learn words and other linguistic forms requires active employment of a theory of mind with which to interpret the intentions behind the speech that they hear (Bloom 2000). The ability to follow the focus of another person's mind, for example by following their eyes or pointing gestures, seems to be essential to normal language acquisition. These two sets of findings and arguments seem to conflict sharply with one another. I will try to show that the evidence behind these opposing positions can be reconciled if we adopt a certain view of language perception and a flexible enough view of the different ways in which one mind can recognize another.

The studies that show that the development of a 'theory of mind' depends on the prior development of rather sophisticated linguistic skills have relied almost exclusively on what are termed 'false-belief tests' to determine when their subjects have acquired a theory of mind. These tests require the subject to recognize a situation in which

embedding. It is true that Grice himself apparently had in mind purely behavioral criteria in the spirit of Gilbert Ryle (and Daniel Dennett 1987), but that alternative is surely quite unattractive in a context that uses Gricean intentions in an explanatory way. Mere behaviors don't explain anything. Only their underlying causes are explanatory.

<sup>7</sup> For a review of the literature and an excellent discussion see Garfield et al. (2001) (which strikes me as a definitive essay on this subject).

misleading circumstances would naturally give rise to a false belief, or they require the subject to recognize that one of their own prior beliefs has been false. What these tests seem to probe is the point at which subjects begin to recognize both their own thoughts and the thoughts of others as taking a form analogous to that of intentional representations such as sentences; the mark of an intentional representation is, of course, that, unlike natural signs, it can be false. As emphasized particularly by Wilfred Sellars,<sup>8</sup> mental states are a bit like theoretical entities. They are not readily publicly observable, nor does the ability to have thoughts imply the ability to think about one's thoughts, any more than the ability of a baby to experience pains and itches implies that the baby has concepts of pains and itches. The idea that there are such things as thoughts which are somewhat like sentences in being sometimes true and sometimes false is, indeed, a sort of 'theory' of mind, call it a 'representational theory of mind'. To acquire a representational theory of mind is a considerable achievement, quite probably a historical cultural achievement, and certainly a considerable achievement for the individual. There is nothing surprising, then, in the fact that there is some delay even in normal children in acquiring this sort of theory of mind. The puzzle is to understand how very young children can be aware of the intentions and of the focus of attention of those from whom they learn language without yet having this sort of sophisticated theory of mind.

I will make three suggestions which when taken together may explain how this is possible. The first is a fairly simple suggestion about how the purposiveness of another's actions may be understood and taken into account without employing a representational theory of mind. The second is a more radical suggestion about the psychological processing involved in language interpretation. I argue that interpreting the meaning of what you hear through the medium of speech sounds that impinge on your ears is much like interpreting the meaning of what you see through the medium of light patterns that impinge on your eyes. When communication proceeds 'Normally'<sup>9</sup>—that is, in the

<sup>8</sup> Best known is Sellars (1956).

<sup>9</sup> The capitalization is to distinguish this sense of 'normal' from one that implies average or most usual. See my (1984: chs. 1–2).



way that accounts for the ultimate survival of linguistic forms and of stable speaker and hearer dispositions in their use—it is the *world*, not meanings, and not speaker intentions, that is immediately perceived when language is understood. Natural language operates on the senses in very much the same way that natural signs such as nonspeech sounds or light operate on them. Both require to be interpreted in context, both require a great deal of filling in or ‘enrichment’, and both sometimes produce illusion. Third, I will suggest how indications of where the speaker’s attention is focused can be recognized during the process of language interpretation without the hearer’s having to employ a representation of the speaker’s mind or of its contents. Together these suggestions are designed to explain how a hearer can come to understand, believe, or do what a speaker intends without having to think about that intention as such. The hearer needs to understand what the speaker intends, but not under that description.

First, then, about recognizing purposes. I began this chapter by mentioning a number of forms that purposes can take other than the form of explicit intentions. Compatibly, recognizing purpose in an activity need not be the same thing as projecting an explicit intention as its cause. Many behaviors of lower animals that we do not take to result from explicit intentions are nonetheless quite obviously purposive. Moreover, it is clear that many animals, perhaps most mammals, are capable of recognizing one another’s purposes. Dogs and cats, for example, could neither play together nor fight with each other without having some grasp of the directedness of one another’s behaviors so as to anticipate one another’s moves. A simple form that this can take is understanding purposive behavior to be ‘goal-directed’; that is, as part of a flexible pattern with a strong tendency to produce a given effect regardless of interfering circumstances. The animal that is goal-directed behaves so as to reach a certain result one way or another despite potentially deflecting forces. This has to be the understanding of both the dog and the squirrel it is chasing, and of the cat as it plays or fights with another cat. Each has to understand the other as being attracted toward certain goals in this manner, as being strongly disposed to continue a particular spurt of behavior until a certain kind of result is reached. Goals appear in this light as though they were future states that are

causing present events to happen. Aristotle called purposes 'final causes', a term that expresses this idea rather neatly.

Plausibly, very small children understand the purposiveness of other people in exactly this sort of way. Studies that show, for example, that a small child can identify the candy on which another child's eyes are focused and will volunteer that that is the one that the other child wants need not indicate possession of a representational theory of mind. The child understands toward which candy bar the other child's behavior will be directed. Similarly, children lacking a representational theory of mind need have no problem understanding that someone else is trying to help them. People who cooperate are directed toward the same goals. The young child often finds that its mother is directed towards the child's own goals and easily learns ways to bring this felicitous circumstance about. Similarly, no representational theory of mind is needed to understand that the speech of another is purposeful, to attend to its purpose, or to attempt to divine its purpose.

But this doesn't answer the deeper question of what children understand the purpose of another person's speech to be, or how they divine what this purpose is in specific instances. To answer that question I must first expand on the second theme mentioned above, the proposal on language understanding. This proposal was defended at some length in my (2000, 2004) and is discussed above, at the end of Chapter 6. I will briefly review my conclusions and then add some clarifications important for the present issue.

The basic idea is that during Normal conversation it is not language that is most directly perceived by the hearer but rather the world that is most directly perceived *through* language. Distal states of affairs are perceived through speech sounds just as they may be perceived, for example, through the medium of structured light during normal vision. But a number of prejudicial assumptions have to be set aside to grasp this as a genuine possibility. One assumption is that believing what one hears or sees directly must be crucially different from believing what one hears said in that the former is (1) so much more reliable and (2) alone yields information about the spatial and temporal relation of the perceiver to what is perceived. I have argued that these are differences of degree only and not relevant to the basic similarities germane

to an understanding of how language is normally processed (2000: ch. 6; 2004; Ch. 9 above). A second assumption is that acquiring beliefs through the testimony of language must involve inference from premises about the speech sounds heard, hence must be indirect in a way that ordinary perception is not. In (2004: ch. 9) I discussed the distinction between direct and indirect perception, in particular the neurology of these, and concluded that there is no reason to suppose that one of these ways of forming beliefs is more direct than the other. A related assumption is that every end organ of sense has a proprietary level, a single depth, at which it perceives directly, so that if it is possible to perceive sentences directly, it cannot also be possible to perceive their distal causes in the world directly. I will say some further words about this third assumption below. Next is the assumption that the objects of perception are always perfectly concrete, never abstract, whereas the information given in language is always highly abstract and sketchy. Last is the claim that because the information provided by language is so abstract and sketchy, the only way for a hearer to fill out this sketchy information to yield the richer understandings we routinely communicate in conversation is through inference about the probable intentions of speakers, a kind of inference obviously not required for ordinary perception. These last two assumptions will be discussed here for the first time. I will argue that ordinary perception is also abstract. On the other hand, I will argue that much that is presented in ordinary perception is not given in the actual data that informs perception but is 'filled in', as it were, during perceptual processing; nor, obviously, does this filling in involve considering anyone's intentions. The same, I will argue, applies to perception of the world through language.

I have used the different levels at which a television screen can mediate perception as an example of shifting levels of direct perception. Usually, one sees at the level of the objects and events depicted; one sees the newscaster frown or smile or sees him point to the map. At other times one might study the dots on the screen, noticing patterns in the static. The TV repairman may be able to see through the dots to perceive exactly what part of the mechanism inside is ailing.<sup>10</sup> The housewife

<sup>10</sup> This perception may be just as direct as the auditory perception of the auto mechanic that the trouble is your main bearing is shot, or the perception of the physician who is able to read an echogram.

sees the grime on the screen and reaches for a rag to clean it. Looking out of a train window when the inside of the carriage is well lit, sometimes one can either watch the people reflected inside or watch the scenery outside with equal ease, neither vision interfering with the other. Similarly, although one normally hears through the language one hears to the events being described, educated adults can attend to phonemes instead if they try, or to phonetics (say, in order to place a foreign accent), or to words, to morphs, to sentences, or to grammatical forms. That words and sentences are not what we *usually* hear (any more than light is what we usually see) is suggested by the ease with which people can remember the content of what was said ten seconds ago in contrast to the difficulty they have in remembering the words themselves. Little children usually cannot segment words into their component phonemes or recognize phonemes as recurring entities at all until they are 5 or 6 (Lieberman et al. 1974). They also have difficulty with the very concept of a word.<sup>11</sup>

Indeed, I think there is evidence that small children don't perceive language at the level of words and sentences for cognitive purposes at all. Consider, for example, the following dialogue (usually used to illustrate that small children are immune to attempts to correct their speech, hence that they don't learn to talk correctly through the imposition of social sanctions):

Child: Nobody don't like me.

Mother: No, say 'Nobody likes me'.

Child: Nobody don't like me.

[Eight repetitions of this dialogue.]

Mother: No, now listen carefully; say '*Nobody likes me*'.

Child: Oh! Nobody don't likes me. (McNeill 1966)

What seems pretty obvious here is that the child simply is not hearing the words, any more than you usually see the dots on the television screen. Both of you are bypassing the vehicle, focusing your powers of perception directly on the content. Compatibly, three- to four-year-olds are typically unable to say whether they have just observed something for themselves or whether they have just been told about it, mostly

<sup>11</sup> Susan Carey, private correspondence.

saying that they have observed what they really have only been told.<sup>12</sup> Several possible interpretations for this have been offered, none of which seem wholly adequate. No one has suggested, however, that for the child, perception through language might just be perception, observing for oneself.<sup>13</sup> Compare observing for oneself through a telescope.

That perception is often abstract<sup>14</sup> is illustrated by occasions on which one can easily perceive the presence of a determinable property without perceiving its determinate value. For example, you can see that an object is some shade or other of red but would need a better light to tell which shade. Or you may clearly perceive that something is looming or moving over there in the dark or that something just shot by you, but not be able to perceive what it is. More striking testimony to the abstract nature of perception are certain illusions of motion; for example, the waterfall illusion (Crane 1989) which presents to the perceiver a landscape that is at once clearly moving and clearly staying in exactly the same place. It moves, but not from one place to another. Contradictions, presumably, are not found in completely concrete representations but only in abstract ones. The world is presented through language very abstractly, but this does not argue against the essentially perceptual nature of the presenting.

Perception may be abstract, but, on the other hand, it is generally more concrete than dictated by the actual data available to the perceptual process. What is perceived is filled in by context. One perceives a chair or a cat, but what is responsible for one's perception is only one side of a part of the chair or the cat. An animal nearly completely occluded behind bushes or any surface with some small openings is

<sup>12</sup> Gopnik and Graf (1988); O'Neill and Gopnik (1991); Whitcombe and Robinson (2000); Wimmer, Hogrefe, and Perner (1988).

<sup>13</sup> Robyn Carston writes that there is some evidence that young children treat what they have been told as being less reliable than what they observe for themselves (private correspondence). As Carston notes, this is not inconsistent with their failure to perceive linguistic forms as such. It may be that children, or that we all, also favor what we have seen over what we have felt or over what we have perceived through the medium of nonlinguistic sounds.

<sup>14</sup> Indeed, in another context, I would argue that it is intrinsically abstract.

easily seen for what it is if its color contrasts with the foreground or if it is moving. If a surrogate dime that is distinctly larger than a real dime is placed at a distance in front of you, in the absence of clear depth clues you will perceive the dime as being dime-sized but closer than it actually is. That is, size is perceptually filled in according to context rather than actually seen, and so is depth. If you listen to the syllable 'ga' while watching the apparent speaker saying 'ba' you will compellingly *hear* the syllable 'da' (McGurk and MacDonald 1976). Foregrounded objects and their backgrounds are reported more accurately when they are consistent with common experience than objects in inconsistent settings (Davenport and Potter 2004). If tiny lights are attached to a person's main joints and a video made showing only these points of light as the person walks about in the dark, watching the video one distinctly sees a person walking about (Johansson 1975). A circle with two marks above inside and a line below is seen as a face, the angles of the marks and curvature of the line determining whether the face is happy, angry, or sad. Stick figures are readily seen as people or as animals, standing, walking, running, lifting things, boxing, waving goodbye. If you see one of Wittgenstein's duck-rabbits surrounded by drawings of ducks you will see it as a duck, but if surrounded by rabbits then as a rabbit. If you draw in water surrounding it, you will again see it as a duck. Consider how you immediately see a certain squiggly line in a cartoon strip to be shoes in the character's hand rather than, say, a bottle because the fellow is obviously tiptoeing up the stairs in stockings, the clock on the wall saying 2.30. The examples go on and on.

Considering this, if language is a medium of direct perception, it certainly is not surprising that half a sentence, or even a single word uttered in the right context, often can convey the same as a full sentence would have. It is not surprising, for example, that when some one calls that they are ready, one generally knows for what they are ready. Nor is it surprising that what is conveyed by language can be much more concrete than the abstract or indeterminate semantic meaning carried by linguistic convention. For example, when one hears that Jane is sad about losing her watch and Jenny sad about losing her husband, one understands these to be rather different kinds of unhappiness, and when one hears that France is hexagonal, that a building is hexagonal,

and that a machine nut is hexagonal one interprets these hexagonals to have different degrees of perfection. The semantics of a possessive allows it to refer to any pairing relation coupling possessors uniquely with things possessed, such as ownership, physical possession, current responsibility for, and so forth, so that 'John's book' may be the one he owns or carries, the one he wrote or bought or brought, and so forth. Still it is not surprising that, in context, it usually is easy to perceive immediately what specific kind of relation is being represented. As for metaphors, gross distortions of the right kind are sometimes more readily recognized in perception than the same items shown more accurately, as illustrated by good caricatures of people's faces. Metaphor is similar. It is often understood effortlessly in context, even sometimes striking a somewhat humorous note just as pictorial caricatures do. I'm certainly not suggesting that it is obvious how perceptual identifications of these various kinds are achieved! But it does seem clear that no thoughts of speaker's minds need be involved in the process.

Most linguistic forms that have multiple senses are effortlessly disambiguated in linguistic context. Other forms are disambiguated just as easily given the immediate external context. One only has to know that we are playing blackjack ('Hit me!'), or that we are eating together ('Can you pass the butter?'). Often, perhaps typically, no consideration of speaker purposes is needed for these interpretations any more than, say, when interpreting what stick figures are doing or perceiving that it's a squirrel rather than a rabbit (obviously) that just ran up the tree. But, of course, there are also many occasions when immediate context is not enough. Instead, an understanding of what the speaker is focusing on is needed both for interpreting what is being said conventionally and for interpreting what may be being meant beyond the conventional.

If someone tells you 'John sprained his ankle' when no John is present, despite the fact that this token of 'John' is being used in an entirely conventional way and has a proprietary conventional referent, clearly you will be lost, it seems, unless you know something about the speaker. The *phonologically*-typed word 'John' is a hugely multiple homonym that divides into thousands of semantic least types, each conventionally referring to a different person, and you need somehow to discern which

of these various lineages this particular token of 'John' has been copied from. Similarly, the content of quantifiers and definite descriptions is nearly always determined by the domain the speaker is focusing on, not just by the immediate context of speaking. And of course there are numerous other examples. These are the cases that have been thought most clearly to require that the hearer think about the speaker's intentions.

But if thinking about speaker intentions is required, does it follow that small children don't ever understand such references? On the other hand, if it is not required, should we suppose that adults employ a representational theory of mind where children are able to do without? The trick is to show how it is possible to make use of natural information about where another person's mind is focused in interpreting their words without employing a representational theory of mind. If we can understand how that is possible, we will have a hold on three subjects of interest. We will have a hold on how children manage to learn language so fast and so accurately (Bloom 2000). We will have a hold on how hearers manage to interpret many nonconventional uses of language; for presumably they do so in much the same way that children manage in first instances to interpret conventional uses. And we will have a hold on how references that can't be interpreted from linguistic form and immediate context alone but only by following the focus of the speaker's mind might be understood both by children and adults without relying on a representational theory of mind. Employing a fully-fledged representational theory of mind in order to interpret what another person is saying or meaning is always a possibility for an adult, but it is plausible, I believe, that this mechanism is seldom required.

We have seen how small children can recognize what another person is trying to do by understanding their activity as goal-directed. Now take seriously that understanding language is originally just a way of perceiving the world through another medium. As the infant matures she learns what things feel like, what those same things look like, what they sound like, often what they smell and taste like, and she learns what they sound like as filtered through the perception and speech of the people around her. She understands speakers as purposefully *showing* her things when they speak, and she tries to interpret what she is being



shown. But how can she tell which aspects of the world around her are being manifested through which speech sounds? How can she make out what it is she is hearing about?

In the case of nonspeech sounds, she can tell, in large part, because hearing is directional. She has learned or innately knows to look or feel in the place or direction from which the sound comes. In this way she soon learns to perceive many ordinary sounds as bearing information about ordinary happenings, such as doors closing, drawers opening, spoons scraping, the eggbeater beating, the toaster popping up, the toy that she drops hitting the floor, her own snuffles and cries, her rustling clothes, and her footsteps. We could put it the other way around just as well. She learns to understand what it is she is *seeing* by its direction relative to what she hears and feels. The capacity to identify the same as the same from any of numerous perspectives including through various sensory modalities is of the essence of conceptualization (Millikan 2000).

Now consider what the child hears through language. When things go Normally, that is in the way that accounts for the survival of languages, another person who tells you things is like a pair of binoculars, or a camera, or a television set. They pick up patterns in the ambient energy surrounding them, interpret these patterns as manifesting certain configurations of distal objects, translate this interpreted structure into new patterns of structured energy, and pipe them over to you. Just as the child's brain (like that of a chimpanzee) is equipped at birth with a neuronal organization that is easily tuned to interpret the visual arrays around it, the child's brain (*unlike* that of a chimpanzee) is equipped with a neuronal organization that is easily tuned to interpret the kinds of informational patterns that language presents. You know what you are seeing through binoculars in part because you know in which direction the binoculars are pointing. Similarly, whether this is learned or innate, children grasp that what people talking to them are talking about is often where these people's eyes or hands are pointing. Nor do they have to think about the construction of the insides of these people's minds in order to understand this, any more than one has to think about how the binoculars are constructed in order to tell what one is seeing through them. Thus the phenomenon of joint looking and joint attention, so much discussed in the literature.

But, of course, people also talk about many things they are not observing. Consider then how you recognize what's in a photograph. Some things in the photograph you may recognize straight off, perhaps certain people, or a very familiar configuration of buildings. Similarly, some of the words talking about things not currently present will already be familiar to the child. At least two other techniques for recognition may be available as well.

The photograph shows a real configuration of various objects and properties on which a camera was once focused. Because the scene was once real, recognizing what is in one part of the picture may give context for recognizing what's in another part. If that's a tiger, then that's bars in the zoo and, now it's clear, that's another cage with a black panther in it, not a toy or a domestic cat. That's Aunt Sadie in the foreground so it must be her new house in the background and the baby pulling the cap down over its eyes must be Billy. Similarly, when part of what is being talked about is understood, what past event or what general kind of event the speaker is focusing on may become clear, bringing that event into focus for the listening child as well. In this way, the content of various new words may be recognized and remembered. Somewhat similarly, if you are telling me about an event or a kind of event that I recognize, my mind will be focused where your mind is focused, and I will understand the proper names you use, your descriptions, and the domains of your quantifiers accordingly, without any concern for what's inside your mind. We are focusing on the same scene or the same type of scene.

Also, one may know where the camera that took a certain photograph has been shooting. If these pictures came from your camera just home from Italy, then I know that would not be Notre Dame but may be Milan Cathedral. Or since the picture of the elderly couple falls on the roll between the picture of Aunt Sadie and the one of baby Billy in the bath, the couple must be her neighbors who came over for tea. Similarly, people talk about what they have observed in the past, most often about what they have recently observed. If my husband announces as he returns from school that 'John' has sprained his ankle I will interpret him in one way, but in other ways if he announces it as he returns from playing tennis or from playing music with friends. What

people talk about depends on where they have been, what they have heard about or experienced. Thus the expression, after having difficulty understanding what someone is talking about, 'I had no idea where he was coming from!' Most frequently, people talk about what they have just experienced. As a child can understand, say, that what a digital camera shows on its monitor will be something to which it has been exposed, children can understand that people are often talking about what they have experienced. Mother is talking to Daddy about what they, the child and Mother, just did in the park. And children can do this without wielding a representational theory of mind, just as they don't need a theory of how the digital camera works inside.

Earlier, I mentioned that the child learns to understand what she is seeing by finding out how it sounds, how it feels, and perhaps how it tastes and smells exactly as she learns to understand what she hears, feels, or tastes by finding out how it looks. No sensory modality exhibits the true nature of a sensed object more than does any other, nor are the natural signs that any one sensory modality relies on more arbitrary, from the standpoint of the interpreter, than for another. The look of a thing no more exhibits its one true essence than does its feel or its smell. The same holds true for perception through language. Objects that a child has only perceived through language, objects that she has heard talked about but has never encountered more directly, may be thought of and judged about as easily and directly as objects she has seen but not heard named. Knowing the meaning of a word in linguistic context is exactly like knowing the meaning of a sight or of a sound in context, nor is the meaning of one more arbitrary than that of the other. The word allows you to recognize the object you are receiving information about exactly as would its sight or its sound. The meaning of a sight or a sound is learned or amplified for one when one learns the name for its object/referent exactly as the meaning of a word is learned or amplified when one learns how its object/referent looks or feels. Once initiated into the realm of language, language begins to speak to the child directly, just as sight, sound, and touch do. 'The child begins to perceive the world not only through his eyes, but also through his speech . . . speech becomes an essential part of his cognitive development' (Vygotksy 1978: 32).

I have said that following the mental focus of another person can be like following the focus of binoculars or of a camera. On the other hand, so as to be understood reliably, speakers also learn to leave trail markers as they shift their focus from one domain to another. They accommodate their discourse to the hearer, taking account of what the hearer might be expected already to know or to have experienced. Whether leaving trails tailored to accommodate the individual hearer requires the speaker to employ a representational theory of mind might be a good question for empirical study. For example, how good are small children at tailoring their conversation to the needs of the particular hearer, and if they are good at it, how do they do it? But I'm not going to speculate about that.

Instead, I will finish with some comments about Gricean implicatures, for it is producing and understanding Gricean implicatures that has traditionally been taken most obviously to require thoughts about thoughts. I have suggested that many interpretations of language forms that Grice described as generalized implicatures may actually be entirely conventional. For example, if there are forms that any cooperative speaker will use only in certain circumstances, and if language conventions proliferate because they are serving cooperative functions, then these forms will inevitably come to mean that those certain circumstances are realized. Thus, if the form 'Some As are B' will be used by a cooperative speaker only when she does not possess information that all of the As are B, then in circumstances in which the speaker may be presumed to know whether all As are B, 'Some As are B' will actually entail that not all As are B. Nor does a hearer need to employ a representational theory of mind to be sensitive, in many cases, to whether the speaker has the relevant information. Knowing that another has certain information can involve no more than knowing the other has been exposed to this information or anticipating that the other could exhibit this information through language.

Ockham's razor cannot discipline the ways that people actually process forms, preventing them from purposefully reproducing and routinely reading context as part of the sign to be interpreted. But if there are forms such as 'Some As are B' that have different readings in different contexts, then out of context they are polysemic or ambiguous.

And because it is perfectly obvious that words are part of the communication system but not so obvious, in many cases, that the context in which words are uttered (or the tone of voice in which they are uttered) is part of that system, it is easy for speakers who are not entirely ingenuous to claim that they have not 'said' what they have conventionally signified because their words (taken alone) could have meant something different. Similarly, children will sometimes claim in fun to have 'said' only what a certain form would mean if read compositionally, whereas the form taken as a whole conventionally has a different or added meaning: 'I didn't ask her for more cookies; I just asked whether she could reach them!'

It is true, however, that observations of this kind don't apply to many of the uses of language that Grice would have labeled 'particularized implicatures'. If A says to B 'Let's go play tennis' and B replies 'It's raining', the obviously intended interpretation that B is unwilling to play tennis surely is not conveyed conventionally. On the other hand, without entertaining any thoughts of what's in B's mind, the mere information that it is raining would undoubtedly be enough to squelch A's hope that B would be disposed to play tennis, hence would serve as a reply to A's invitation. And that B is purposefully replying to A's invitation—that B is goal-directed toward revealing to A his lack of inclination (read that behaviorally) toward playing tennis—can also be understood without employing a representational theory of mind. In general, when a speaker's purpose in what he says involves bringing something closely associated or implied to a hearer's mind, a hearer can understand that this effect on her is purposive without using a theory of mind. She can do this just as she can understand that the speaker is purposefully insulting, hurting, or instead trying to help or comfort with his words. I certainly don't want to claim that all particularized Gricean implicatures can be handled in this sort of way. But how far it is possible to get in understanding such implicatures prior to acquiring a representational theory of mind is not, I believe, obvious a priori. It remains an interesting question for empirical study.

It is, of course, obvious that for a hearer to understand what a speaker means is for her to grasp what the speaker intends (or at least what he purposes, taking into account that not all purposes take the form of

explicit intentions). The question that has been at issue here is how frequently it is necessary that the hearer grasp 'what the speaker intends' where that phrase is to be read *opaquely*, and how frequently as read merely transparently. If the hearer grasps that which the speaker is trying to impart, that is if she grasps that such-and-such is the case or that she is to do such-and-such, that will usually be enough. In deriving this, the hearer need not represent to herself that the speaker has imparting this as his intention. Indeed, in the usual case, for the hearer to diagnose the hearer's intentions she must first interpret the speaker's words (in context), not the other way around.

## REFERENCES

- Bloom, P. (2000), *How Children Learn the Meanings of Words* (Cambridge, Mass.: MIT Press).
- Crane, T. (1989), 'The Waterfall Illusion', *Analysis*, 483: 142–7.
- Davenport, J. L., and Potter, M. C. (2004), 'Scene Consistency in Object and Background Perception', *Psychological Science*, 15/8: 559–64.
- Davis, S. (1991) (ed.), *Pragmatics: A Reader* (Oxford: Oxford University Press).
- Dawkins, R. (1976, rev. 1989), *The Selfish Gene* (Oxford: Oxford University Press).
- Dennett, D. C. (1987), 'Intentional Systems in Cognitive Ethology: The "Panglossian Paradigm" Defended', in *The Intentional Stance* (Cambridge, Mass.: MIT Press), 237–69.
- Garfield, J., Peterson, C., and Perry, T. (2001), 'Social Cognition, Language Acquisition and the Development of the Theory of Mind', *Mind and Language*, 16/5: 494–541.
- Gopnik, A., and Graf, P. (1988), 'Knowing How You Know: Young Children's Ability to Remember and Identify the Sources of their Beliefs', *Child Development*, 56: 1366–71.
- Harnish, R. M. (1976), 'Logical Form and Implicature', in T. D. Bever et al. (eds.), *An Integrated Theory of Linguistic Ability* (New York: Crowell); repr. in Davis (1991: 316–54).
- Johansson, G. (1975), 'Visual Motion Perception', *Scientific American*, 232: 76–89.
- Lieberman, A. M. (1999), 'The Reading Researcher and the Reading Teacher Need the Right Theory of Speech', *Scientific Studies of Reading*, 3: 95–111.

- Lieberman, I. Y., Shankweiler, D., Fisher, F. W., and Carter, B. (1974), 'Explicit Syllable and Phoneme Segmentation in the Young Child', *Journal of Experimental Child Psychology*, 18: 201–12.
- McGurk, H., and MacDonald, T. (1976), 'Hearing Lips and Seeing Voices', *Nature*, 264: 746–8.
- McNeill, D. (1966), 'Developmental Psycholinguistics', in F. Smith and G. A. Miller (eds.), *The Genesis of Language* (Cambridge, Mass.: MIT Press).
- Millikan, R. G. (1984), *Language, Thought, and Other Biological Categories* (Cambridge, Mass.: MIT Press).
- (2000), *On Clear and Confused Ideas* (Cambridge: Cambridge University Press).
- (2002), '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).
- (2004), *Varieties of Meaning* (Cambridge, Mass.: MIT Press).
- Morgan, J. L. (1978), 'Two Types of Convention in Indirect Speech Acts', in P. Cole (ed.), *Syntax and Semantics, ix. Pragmatics* (San Diego, Calif.: Academic Press), 261–80; repr. in Davis (1991: 242–53).
- O'Neill, D. K., and Gopnik, A. (1991), 'Young Children's Ability to Identify the Sources of their Beliefs', *Developmental Psychology*, 27: 390–7.
- Searle, J. (1975), 'Indirect Speech Acts', in P. Cole (ed.), *Syntax and Semantics, ix. Pragmatics* (San Diego Calif.: Academic Press), 59–82; repr. in Davis (1991: 265–77).
- Sellars, W. (1956), 'Empiricism and the Philosophy of Mind', in K. Gunderson (ed.), *Minnesota Studies in the Philosophy of Science*, 1. (Minneapolis, Minn.: University of Minnesota Press), 253–329; repr. in Sellars (1963).
- (1963), *Science, Perception and Reality* (London: Routledge and Kegan Paul; repr. Atascadero, Calif.: Ridgeview, 1991).
- Vygotsky, L. (1978), *Mind in Society: The Development of Higher Psychological Processes*, ed. M. Cole, V. Johns-Steiner, A. Cribner, and E. Souberman (Cambridge, Mass.: Harvard University Press).
- Wheeler, S. (1972), 'Attributives and Their Modifiers', *Nous*, 6/4: 310–34.
- Whitcombe, E. L., and Robinson, E. J. (2000), 'Children's Decisions about What to Believe and their Ability to Report the Source of their Belief', *Cognitive Development*, 15/3: 329–46.
- Wimmer, H., Hogrefe, G. J., and Perner, J. (1988), 'Children's Understanding of Informational Access as a Source of Knowledge', *Child Development*, 59: 386–96.

# Index

- activity  
  conventional 146, 147  
  patterns of 4, 5, 8, 9, 18, 20, 29,  
    33, 40, 146–8, 180  
  purpose in an 206, 213  
  shape of 149
- adaptation 57, 101, 156, 177
- affordances 174
- Alston, W. 121, 122
- ambiguities 18, 202
- Anscombe, G. E. M. 166
- Aristotle 108, 129, 132, 207
- artifacts 109
- attitudes 28, 184  
  intentional 74, 90  
  propositional 173
- Austin, J. L. 19, 59, 62, 139, 140,  
  151, 154, 158, 163
- Bach, K. 11, 19, 27, 56, 139, 145
- bees 64, 83–7, 94, 96–9, 156, 167, 174
- behavior 5, 56, 57, 79, 80, 85, 87, 88, 163,  
  174, 178, 179, 181, 187, 188, 206, 207  
  adaptive 101  
  conventional 3, 4, 7, 8, 15, 56  
  coordinated 177  
  imitative 175  
  learned 93, 95, 156, 187  
  natural 19, 21, 146  
  normal 156  
  patterns of 7, 8  
  prediction 136  
  regularity in 55, 56, 59  
  regulated 20  
  rule obeying 80  
  social 20  
  unconventional 3
- uniform 26, 27
- beliefs 6, 13, 27–9, 45–7, 62, 73, 85, 86,  
  88, 93, 95, 117, 135, 136, 153, 157,  
  169, 170, 172, 176, 190, 203,  
  204, 208  
  about words 65  
  descriptive 184  
  false 204, 205  
  forming 59  
  interpreting 172  
  true 58, 63, 94, 96, 98, 170, 172, 183,  
  190
- Block, N. 129
- Bloom, P. 194, 204, 213
- Bradley, F. H. 126
- Brandon, R. 77, 79, 81, 82, 84, 86, 89
- Brentano, F. 168
- Burge, T. 67, 117, 128, 133
- Carey, S. 65
- causal roles 125
- chameleon 170
- Chomsky, N. 2, 24, 25, 26, 28, 36, 37, 38,  
  39, 49, 50, 119
- chunking 198
- classes 106–11, 143
- cognition 87, 95, 136
- communication 19, 24–8, 39, 43, 46, 50,  
  58, 63, 124, 139, 160, 189, 190, 191,  
  203, 205, 218
- conceptions 54, 69–74  
  public 73
- concepts 65–73, 84, 85, 89, 93, 95, 103,  
  104, 112–16, 119, 121, 124, 125, 131,  
  137, 170, 184, 205  
  learning 104  
  thick 180, 184, 185



- conceptual analysis 66, 67, 121–5, 128,  
129, 132
- conceptual components 54, 69, 70, 71, 74
- conceptual roles 78
- conditioning 48, 79, 85, 187, 188
- conformity 2, 6, 11–16, 41, 56, 59, 87,  
148, 169
- conspicuous 106, 155, 157, 167
- constancies  
perceptual 67, 68, 88
- content 90, 126, 140, 173, 209, 213, 215  
propositional 58, 190  
representational 171, 173  
thought 123, 129, 135
- conventions 1–22, 25–50, 55–9, 63, 79,  
86, 139, 140, 142–55, 158, 159, 162,  
163, 164, 182, 189, 190, 192, 194,  
200, 212, 213, 217  
blind 15, 41, 42  
coordination 40, 41  
crisscrossing 18, 25, 38, 60, 202  
grammatical 36  
half-blind 13  
leader–follower 12, 14  
linguistic 1, 2, 13, 16, 20, 22, 25–33,  
42, 43, 44, 46, 48, 54, 55–61, 100,  
141, 145, 152, 169, 188, 191, 192,  
195, 201, 202, 211, 217  
natural 2, 3, 9, 18, 20, 21  
open 42  
public 37, 50  
regulated 162  
social 55, 142
- cooperation 12, 42, 53, 54, 63, 85, 86,  
87, 90, 96, 152, 155, 157–61, 190,  
192, 217
- coordinations 2, 9, 10, 11, 12, 39–48, 55,  
56, 59, 177, 178, 191  
beliefs in 9  
blind 11, 12, 14  
conventional 13, 14, 40, 42, 48, 56, 59  
half-blind 11, 13  
leader–follower 12, 13, 19  
open 11, 12, 42  
social 177
- copying 4–8, 17, 30, 34, 35, 61, 92, 109,  
111, 146, 162, 169, 183, 187, 188,  
192, 193–202, 213
- correspondence 87, 97, 98, 99
- Davidson, D. 19, 49, 50, 74
- Dawkins, R. 39, 187, 189
- demonstratives 2, 66, 200
- descriptions 59, 69, 71, 73, 112, 116, 117,  
123, 128, 129, 132, 133, 147, 154,  
175, 181, 185, 206, 215  
definite 62, 65, 66, 74, 124, 126, 213  
empty 123  
identifying 116  
individuating 112, 113
- descriptives 63, 88, 90, 166, 167, 172,  
173, 175, 176, 178, 181, 183, 184
- determinables 210
- directives 16, 59, 63, 88, 90, 99, 152, 166,  
167, 173, 175, 176, 178, 181, 183, 184
- dispositions 46, 54, 64, 71, 81, 82, 83, 93,  
95, 101, 104, 110, 113, 124, 126, 129,  
157, 178, 184, 206
- Donnellan, K. 62, 65, 75
- Dretske, F. 129
- Dummett, M. 25
- Eldredge, N. 111
- epistemology 72, 103
- essences 110, 128, 178
- Evans, G. 67
- evolution 19, 36, 43, 46, 61, 83, 94, 102,  
168, 176, 177
- evolutionarily stable solution 191
- existence 14, 15, 26, 39, 73, 92, 98, 135,  
155
- extension 127
- externalism 24, 71, 72, 123, 126,  
129, 130

- fact icons 99  
 fidelity 189  
 final causes 207  
 Fodor, J. 102, 129  
 Foot, P. 184  
 force 142, 143  
     illocutionary 62, 143  
     perlocutionary 62  
 Frege, G. 54, 66, 67, 72, 73, 123, 127, 132  
 functions 2, 7, 8, 26, 90  
     alignment 102  
     biological 39, 40, 53, 93, 94, 95, 109,  
         131, 156, 168, 170, 178, 188  
     changing 196  
     conflicting 160  
     conventional 28, 30, 31, 40, 42, 43, 44,  
         45, 49, 50, 58, 90, 146, 179  
     cooperative 54, 63, 157, 160, 217  
     coordinating 44, 61, 152, 177, 179, 188  
     coordination 40  
     derived 191  
     descriptive 166, 172, 179, 182, 184  
     directive 166, 173, 174, 179, 182, 185  
     explanation of 97  
     failed 168  
     grammatical 44, 46, 47, 58, 94, 182,  
         183, 190  
     interpreter 172  
     linguistic 18, 24–8, 39, 43, 45, 47, 48,  
         50, 53, 54, 58, 61, 65, 85, 90, 93, 95,  
         96, 98, 100, 140, 145, 151, 152, 155,  
         159, 169, 170, 179, 184, 190, 193,  
         195, 196, 197  
     memetic 188, 191, 200  
     proper, *see* proper functions  
     semantic 202  
     speaker-derived 100  
     stabilizing 54, 58, 60–6, 73, 74, 75, 90,  
         94, 95, 96, 98, 100, 188, 191, 192  
  
 Gallistel, C. 174  
 games 14, 77, 82, 178  
  
 Ghiselin, M. T. 106, 108, 111  
 Gibbard, A. 177, 184  
 Gibson, J. J. 87, 174  
 Gilbert, M. 5, 10, 11, 27, 56, 117  
 goal icons 99  
 goals 99, 132, 175, 176, 187,  
     188, 206, 207  
 Gould, S. 111  
 grammar 35, 36, 38, 44, 62, 158, 159, 169  
     universal 2, 30, 35, 189  
 Grice, H. P. 19, 20, 48, 49, 61, 85, 95, 139,  
     140, 141, 156, 159, 160, 169, 170,  
     179, 188, 191, 203, 217, 218  
 groups 27, 29, 32, 109  
  
 Harnish, R. M. 11, 19, 27, 56, 140, 145,  
     198  
 hearers 6, 16, 43, 45–9, 53, 57, 58, 61, 63,  
     73, 79, 85, 90, 93, 94, 95, 153, 157,  
     158, 169, 172, 179, 185, 189, 190,  
     192, 193, 196, 203  
 homeostasis 107  
 Hull, D. 106, 108, 111  
 identity 48, 65, 73, 124, 135  
 illocutionary acts 19, 139–45, 151,  
     154–6, 161, 162, 164  
 illusions 118, 210  
 imperatives 6, 47, 58, 94, 96, 141, 152,  
     157, 160, 172, 179, 191, 196  
 implicatures  
     conversational 191, 196  
     Gricean 48, 49, 61, 169, 170,  
         179, 203, 217  
     particularized 218  
 indexicals 66, 124, 128, 129, 200  
 indicatives 6, 14, 44–7, 58, 62, 94, 96, 98,  
     157, 159, 190  
 individualism 133  
 individuals 71, 107–13, 128  
     concepts of 70  
     names of 69  
     thoughts of 126, 127

- induction 108
- inferences 71, 78, 88–90, 111, 124, 125, 133, 167, 170, 172, 173, 179, 201, 208
- information 45–7, 58, 88, 89, 103, 104, 107, 115, 116, 119, 141, 143, 175, 190, 218
- conceptual 65, 69, 73, 134, 135, 217
- conveying 6, 13, 28, 44, 46
- natural 116, 117, 213
- perceptual 118, 207, 214, 216
- sharing 85
- intension 39, 54, 72, 127
- intentional attitudes 46, 48, 72, 74, 87, 90, 93, 96–8, 102, 123
- intentionality 78, 82, 85, 90, 92, 93, 96, 97, 98, 100, 102, 168
- intentions 13, 58, 78, 79, 85, 86, 95, 96, 99, 131, 140, 141, 144, 145, 150–6, 161–3, 167, 170, 176, 179, 180, 184, 187, 188, 206, 219
- concepts of 204, 205
- cooperative 160
- Gricean 20, 95, 140, 156, 159
- hearer 47
- illocutionary 142, 155
- speaker 22, 28, 100, 141–4, 151–5, 158–60, 163, 164, 169, 201, 203, 206, 208, 213, 219
- internal realism 130
- interrogatives 47, 62, 94, 157, 196
- isomorphism 77, 78, 97, 102
- Jeannerod, M. 175
- judgments 71, 78, 88, 89, 103, 125, 126, 134
- Kant, I. 77
- kinds 71, 101, 102, 124, 128
- of conventions 46
- eternal 109, 110, 111, 114
- historical 61, 106, 107, 108, 109, 111, 114
- names of 70, 133
- natural 70, 104, 126, 128, 133
- Kripke, S. 117, 125, 126, 129
- language 18, 44, 49, 50, 70, 78, 82, 84, 89, 92, 99, 104, 130, 167, 169, 178, 189, 190, 191, 201
- animal 2
- child 204, 205
- community 24, 54, 61, 104, 129, 172
- conceiving through 104
- forms 30, 57, 62, 159, 169, 181
- game 89
- I- 28, 37, 39
- intentionality of 96, 98
- interpretation 201, 205
- learning 25, 37, 66, 84, 125, 133, 157, 194, 195, 213
- mechanisms 36
- natural 1, 2, 6, 29, 33, 38, 46, 56, 61, 75, 157, 169, 202
- non-natural 2
- norms 83
- perception through 69, 116, 117, 118, 119, 207, 209–16
- philosophy of 122, 128, 130, 137
- public 22–8, 33, 37–9, 49, 50, 54, 58, 60, 70, 84, 89, 92–6, 99, 123, 125, 133, 140, 151, 153, 155, 159, 169, 170, 172, 180, 184, 188, 189, 196
- purpose of 24, 27
- semantic 80
- of thought 102
- understanding 58, 172, 207
- use 86, 87, 134, 144, 160, 161, 170, 199, 218
- law 15, 20, 148, 150, 161, 163, 164, 181, 201
- of effect 83
- natural 68
- learning 70, 79, 83, 84, 85, 101, 153, 157
- least types, *see* types, least

- Lewis, D. 1, 2, 5, 7, 9, 10, 11, 16, 26, 27, 29, 33, 40, 41, 55, 56, 59
- Liberman, A. 36
- linguistic idealism 124
- Lockwood, M. 135
- meaning 49, 96, 98, 136, 193  
 changes in 192, 193, 196  
 conception 70, 72  
 conventional 54, 55, 142, 143, 188, 195, 196, 199, 200  
 criterion for 134  
 extensional 123  
 externalism 71, 72, 125, 126, 127, 128, 133  
 holistic 195  
 internalism 66, 130  
 linguistic 37, 49, 53, 54, 58, 63, 89, 92, 122, 124, 125, 132, 133, 134, 197, 205, 216  
 literal 100, 191, 196  
 memetic 202  
 multiple 61  
 public 66, 123, 125, 128, 134  
 referential 123, 124, 125, 130  
 semantic 196, 198, 199, 201, 203, 211  
 senses of 60, 66, 72  
 speaker 49, 54, 55, 189  
 theory of 24  
 vague 198
- Meltzoff, N. 175
- memes 187, 188, 191, 193, 200, 202
- metaphor 17, 29, 44, 47, 49, 61, 192, 193, 195, 196, 212
- Millikan, R. 39, 47, 59, 62, 64, 65, 67, 68, 72, 73, 74, 77, 93, 94, 95, 98, 99, 100, 102, 103, 104, 106, 117, 131, 132, 134, 139, 166, 168, 169, 172, 214
- mind  
 theory of 189, 204, 205, 207, 213, 216, 217, 218
- misrepresentation 168
- modality 88, 98
- models 3, 31, 35, 64, 183, 202
- Moore, K. 175
- Morgan, J. L. 197
- moves  
 chess 15, 82, 148, 149, 182  
 conventional 20, 21, 149, 150–5, 158, 161–4, 180–3  
 inferential 82  
 linguistic 81  
 names of 181, 182, 183  
 regulated 21, 150, 162
- names 74, 78  
 of classes 108  
 common 73  
 proper 68, 124, 126, 129, 132, 215  
 species 107  
 of speech acts 151
- natural selection 36, 83, 84, 94, 157, 168, 187
- natural signs 118, 205, 206, 216
- normativity 79–82, 86, 177
- norms 27, 44, 79, 80, 81, 83, 84  
 common 177  
 conceptual 81, 82, 84, 86  
 coordinating 178  
 public 25  
 role 177  
 semantic 81, 82  
 social 177, 178, 179
- one-to-one assumption 97, 123, 124, 125, 128, 132, 136, 137
- ontology 72, 103, 128, 136
- outcomes  
 conventional 20, 21, 22, 47, 100, 140, 149, 150–5, 158, 162–4, 181–3  
 regulated 21, 162
- Papineau, D. 130
- patterns 5, 7

- patterns (*cont.*)  
of behavior 178  
conventional 4, 14, 15, 16, 17, 18, 22,  
31, 32, 35, 37, 40, 43, 47, 48, 148, 150  
cooperative 42  
coordination 40, 59  
inferential 90, 133  
of interpretation 153  
linguistic 6, 31, 33, 43, 45, 48,  
60, 81, 83, 193  
perceptual 214  
regulated 21, 161  
speaker–hearer 153, 169, 201  
syntactic 3
- perception 67–9, 87, 88, 100, 102, 112,  
117, 118, 125, 174, 176, 204, 207–13,  
216, 219  
constancies, *see* constancies, perceptual  
direct 69, 117, 208, 211
- perlocutionary acts 19, 59, 100
- photographs 215
- Pinker, S. 44
- possessives 212
- pragmatics 2, 89, 90, 100, 188, 201, 203
- precedent 2, 7, 8, 10, 11, 13, 17, 18, 55,  
56, 60, 86, 90, 188, 192
- productivity 36
- projection 78, 83, 88, 97
- proper functions 22, 45, 90–9, 102, 141,  
156–62, 168–74, 176, 181, 182, 185  
cooperative 155, 156, 158, 159  
derived 93, 100, 156, 157, 160, 188  
direct 39, 100
- proper reactions 169
- properties 30, 66, 68, 70, 71, 102, 104,  
107, 108, 110–17, 124, 126, 127,  
133, 185  
identifying 72  
names of 69  
recognition of 113
- propositions 44, 45, 74, 157
- psychology 39, 43, 80, 101, 102, 136
- purposes 95, 141, 145, 162, 176  
of behavior 156, 188  
biological 85, 95, 131, 132  
conventional 140  
cooperative 12, 160, 161  
implicit 99  
lacking 178  
linguistic 22, 26, 48, 50, 58, 140, 152,  
155, 160, 161, 188, 190, 191  
memeic 200  
natural 84, 85, 187, 189, 196  
recognizing 206  
secondary 43  
shared 9, 40, 43, 55, 157  
speaker 6, 55, 57, 94, 100, 139, 140,  
145, 159, 190, 192, 207, 212, 218  
speech act 151, 156, 158, 162  
of thought 131
- Putnam, H. 67, 117, 125, 126, 128, 130,  
131, 133, 134
- quantifiers 200, 213, 215
- Quine, W. V. O. 78, 124, 125, 127,  
133, 134
- Recanati, F. 27, 56, 140, 145
- recognition 67, 68, 69, 70, 88, 102, 104,  
112–19, 131–6, 204, 215, 216
- reference 65, 66, 115, 117, 123, 124, 127,  
129–36  
demonstrative 200  
direct 123, 125, 129  
public 135
- reference class 27
- regularities 14, 16, 26, 29, 33, 41,  
56, 79
- reidentification 69, 103, 104
- relations 30, 108, 116, 124, 127  
historical 127, 128, 129, 130, 132  
perception of 68  
symbiotic 46

- representations 74, 84, 88–90, 97, 99,  
     103, 130, 166–78, 184, 185, 206  
     contradictions in 210  
     descriptive 171–4  
     directive 171, 172, 174  
     imperative 99, 157  
     indicative 98  
     intentional 205  
     perceptual 87, 102, 174  
     pushmi-pullyu 87, 99, 166, 167, 168,  
         170, 173–85  
 reproduction 2–8, 12–19, 21, 30–41, 45,  
     47, 56, 57, 59, 60, 86, 92, 93, 109,  
     110, 146–9, 152–7, 168, 169, 183,  
     188–97, 200, 201  
     aspects of 30, 35  
     of behavior 7  
     counterpart 4, 5, 6, 8, 10  
     definition of 30  
     handshake 4, 5, 9, 10  
     nuts and bolts 4, 31, 45  
     of words 35  
 rigid designators 126  
 Rizzolatti, G. 175  
 rules  
     of chess 17, 32, 147, 163  
     conditional 41  
     conventional 6, 19, 152, 153  
     conventions and 14, 16, 17, 32,  
         142, 143  
     correspondence 99, 104  
     of etiquette 148  
     of inference 78  
     of language 18, 25, 27, 37, 77, 78,  
         79, 82, 83, 124, 125, 133, 169,  
         191, 192  
     normative 80  
     obeying 80  
     in perception 87  
     semantic 45, 63, 84, 98, 99, 169  
     transformation 80  
     of translation 89  
 Russell, B. 126, 127, 132, 137, 198  
 same saying 74  
 sanctions 8, 15, 16, 81, 148, 149, 152,  
     157, 158, 164, 209  
 satisfaction conditions 49, 54, 63, 90, 93,  
     171, 172, 173, 183, 184  
 Schiffer, S. 11, 19, 27, 56, 139, 145  
 Searle, J. 11, 19, 27, 56, 140, 149  
 seed assumption 123, 124, 125, 126, 129,  
     131, 132, 136, 137  
 selection 85, 189, 190  
     genic 93, 107  
     group 177  
 Sellars, W. 48, 64, 77, 78, 79, 80, 81, 82,  
     83, 84, 86, 87, 90, 125, 205  
 semantic/pragmatic distinction 50, 192,  
     193, 194, 201  
 semantics 44, 45, 46, 47, 50, 60, 84, 89,  
     189, 196, 200, 202, 203  
     assessment 80, 81  
     compositional 199  
     inferential role 77, 78  
     mapping functions 53, 54, 63,  
         64, 66, 188  
     of representation 90  
 sense 35, 54, 72, 194, 196  
     public 135  
 sense data 126  
 sentences 63, 64, 65, 74, 75, 78, 99, 123,  
     182, 189, 194, 211  
     identity 73, 135  
     meaning 81, 98  
 Skinner, B. F. 43  
 Sober, E. 177  
 speakers 43, 45, 47, 48, 53, 54, 57, 58, 61,  
     79, 85, 90, 94, 95, 153, 157, 169, 190,  
     193, 196, 203, 217  
     intentions of 39  
 species 34, 46, 61, 83, 101, 106, 107, 114,  
     157, 169, 174  
     as individuals 111

- speech acts 2, 19, 20, 22, 29, 79, 85, 90,  
100, 139, 140, 143, 145, 159, 160,  
161, 164  
K-I 139–45, 151–64  
K-II 139–45, 161–4
- Sperber, D. 18, 161
- Stalnaker, R. 85
- Strawson, P. F. 19, 48, 65, 73, 135, 139,  
140, 141, 142, 143, 144, 145, 150,  
151, 153, 155, 156, 158, 159, 160,  
161, 162, 164
- substances 108, 110, 111, 114,  
115, 116
- syntax 2, 3, 16, 35, 36, 44, 45, 58, 60, 62,  
63, 64, 94, 156, 157, 169, 172, 179,  
190–3, 196, 202
- Tarski, A. 63
- thought 24, 44, 67, 70, 71, 78–85, 90, 92,  
96, 99, 100, 102, 104, 112, 123,  
126–31, 136, 167, 176–9, 187  
rules of 86
- tokens 17, 33, 60, 61, 129, 132, 196, 197,  
199, 202, 203
- tradition 3, 8, 15, 30, 40, 202
- truth 81, 84  
speaking the 157  
truth conditions 44–7, 50, 53, 54,  
63, 65, 81, 90, 93, 97, 98,  
172, 183  
types  
least 62, 63, 66, 69, 193, 198, 202, 203,  
212
- understanding 54, 66, 68–70, 80, 81, 113,  
116, 118, 123, 141, 142, 160, 172,  
195, 206, 207, 214–18  
social norms 178
- voting 17, 19, 20, 21, 147, 181, 182
- Warnock, G. J. 140
- Wheeler, S. 197
- Williams, B. 184
- Wilson, D. 18, 161, 177
- Wittgenstein, L. 63, 71, 77, 85,  
88, 107, 125, 134, 211
- words 17, 33, 34, 35, 104, 133,  
199, 211, 216  
learning 119  
meaning of 123  
recognition through 119  
tokening 35  
typing 34