Of these feature types, to my knowledge only the modality and descriptive features have been investigated in neural terms.

Within the context of our psychological concerns, it seems to me that this overall organization of feature types is not something that can be ascribed to learning. This is a bare skeleton of the framework in terms of which learning can take place. On the other hand, since this organization plays a role in perception, imagery, and action as well as language, we do not have to ascribe it to narrow Universal Grammar: it is a central part of all cognition.

10.7 Application to theory of reference

Section 10.5 showed how the act of creating or understanding an expression with deictic reference can be reduced functionally to the association or binding of a linguistic entity with a percept. We are now in a position to sharpen and generalize this account.

The indexical feature of a percept is the crucial feature for linguistic reference. If there is no indexical feature to which descriptive features can be bound, there is nothing to which a deictic or other referring linguistic expression can be attached either: there is no *'that'* there.

We have already seen how this account extends to some of the problematic entities mentioned in section 10.3. Consider again virtual objects such as the square subtended by four dots. Although there is nothing "actually out there," the visual system constructs a percept with all the right features for "something out there." Hence there is something "out there" in the *perceived* world, and we can refer to it. Next consider the unicorn in my dream. There is a remembered entity which has an indexical feature and descriptive features in the visual modality (in particular a one-horned shape). In memory, it is assigned the valuation *internal*, so it is not experienced any more as having been "in the world." Still, because it has an indexical feature, it can be connected to a referential expression.

Returning to the standard literature on reference: Frege (1892) made the distinction between the sense and the reference of an expression by citing the well-known example *The morning star is the evening star*. In his analysis, two senses are attached to the same reference. We can understand Frege's example in present terms as reporting the merger of indexicals associated with different descriptive features, exactly as in our example of the circle that is sometimes red and sometimes blue. So Frege's problem is not a uniquely *linguistic* problem; rather, it lies in a more general theory of how the f-mind keeps track of individuated entities.

The reverse of merger can also take place. For a long time I thought the literary/cultural theorist Allan Bloom and the literary/cultural theorist Harold Bloom were the same person. In present terms, the descriptive features of both were linked to a single indexical. Then one day, believing that (Allan) Bloom had died a couple years back, I was astonished to come across a new book by (Harold) Bloom. The embarrassing realization "Ohhh! There are *two* Blooms!" forced me to split the indexical and divide the (admittedly minimal) descriptive features among them; in particular, only one was dead. (In a realist theory of reference, who was I referring to by *Bloom*, before my epiphany?)

More generally, the indexical features proposed here play a role altogether parallel to the discourse referents in various approaches within formal semantics such as Discourse Representation Theory (Kamp and Reyle 1993), File Change Semantics (Heim 1989), and Dynamic Semantics (Groenendijk et al. 1996). Hence many of the insights of these approaches can be taken over here, with the proviso that the "model" over which reference is defined should be a psychologically motivated one. (We return to these approaches in sections 12.3 and 12.4.)

10.8 Entities other than objects

Let us go a bit more into what a psychologically motivated model would include, by extending our analysis of deictic reference in sentences like (10), *Hey, look at* that*!* The main point to be established is that indexical features do not always identify *individual objects*: these are just one sort of entity that can be identified by an indexical.

For instance, consider (10) uttered in reaction to muddy water pouring out of a pipe. Here the intended reference may well be the *substance* rather than any particular bounded quantity of it. This distinction is reflected in the mass-count distinction in grammar and more generally in aspectual distinctions such as perfective versus imperfective (see Jackendoff 1991 and references there).

The simple deictic *that* can also be used to refer to non-material entities such as sounds and tactile sensations:

(11) a. Goodness! Did you hear *that*?
b. [Dentist:] Did *that* hurt?

Hankamer and Sag (1976) notice that English contains a substantial number of other deictic expressions that can be accompanied by gestures. Jackendoff (1983) points out that they are every bit as referential as the deictic in (10).

(12) a. Pro-PP:
Please put your coat *here* [pointing] and put your hat *there* [pointing].
He went *thataway* [pointing].

b. Pro-VP:
Can you *do that* [pointing]?
Can you *do this* [demonstrating]?

c. that . . . happen:
That [pointing] had better not *happen* around here.

d. Pro-manner adverbial:
You shuffle cards *thus/so/like this/this way* [demonstrating].
I used to walk *that way/like that* [pointing to someone walking in a funny fashion].

e. Pro-measure expression:
The fish that got away was *this/yay* long [demonstrating].
There were about *this many* people at Joe's party too. [gesturing toward the assembled throng]

f. Pro-time-PP:
You may start . . . right . . . *now*! [clapping]

The analysis of (10) suggests parallel analyses for these examples: the deictic expressions refer to entities in the conceptualized world that can be picked out with the aid of the accompanying gesture. But the entities referred in (12) are not *objects*. In (12a), *here* and *there* refer to locations or places; *thataway* refers to a direction. *Do this* and *do that* in (12b) and *that . . . happen* in (12c) refer to actions or events (goings-on). The expressions in (12d) also refer to actions but draw attention to the manner in which the action is performed. The expressions in (12e) refers to sizes and amounts, and *now* in (12f) refers to a time.

In order to accommodate these possibilities for reference, it is useful to introduce a kind of "ur-feature" that classifies the entity being referred to into an *ontological type*. Each of the ontological types—objects, actions, locations, and so forth—has its own characteristic conditions of identity and individuation. It is a task both for natural language semantics and for cognitive psychology/neuroscience to work out the logic and the characteristic perceptual manifestations of each type. Events, for instance, play a role in many versions of formal semantics, following Davidson (1967); some of the psychophysics of visually perceived events is explored by Michotte (1954) and Cutting (1981). Locations and directions have not played much of a role in formal semantics, but are central in Cognitive Grammar (e.g. Langacker 1987; Talmy 1983; Herskovits 1986; Vandeloise 1986) and have figured in recent experimental research, much of which is described in Bloom et al. (1996).

The ontological categories of sounds, tactile sensations, manners, and distances have not to my knowledge received significant attention in the literature. Let us look for a moment at the logic of sounds. Notice that the two sentences in (13) can be used to describe the same situation.

(13) a. There's that noise again!
b. There's another one of those noises!

(13a) describes the situation as the recurrence of the same individual—hearing the noise again is like seeing the sun again the next day. By contrast, (13b) describes the situation as hearing another sound of the same type (we get to types shortly). A similar ambivalence attaches to other temporally dependent ontological types, for example *There's that pain in my leg again* vs. *There's another of those pains in my leg!*; *It's sunrise/Tuesday again* vs. *It's another sunrise/Tuesday.*

Words appear to be classified as a species of sounds (even when written): *'Star' appears twice in that sentence* vs. (slightly less natural) *There are two 'star's in that sentence. Mahler's Second* (example (5a)) leans more strongly in the direction of the (13a) conceptualization: *They played Mahler's Second six times on tour* is much more natural than *They played six Mahler's Seconds on tour* (musicians do, however, say this sort of thing).[18]

Are all the sorts of entities in (11)–(12) "in the world"? They are certainly not like refrigerators—you can't touch them or move them. In fact, it is odd to say they all *exist* in the way that refrigerators exist ("the length of the fish exists"??). Yet (11)–(12) show that we can pick them out of the perceptual field and use them as the referents of deictic expressions. So we must accord them some dignified status in the conceptualized world—the "model" that serves as the basis for linguistic reference.

Now notice what has just happened. Up until this section we have been concerned with reference to objects, and we have used perceptual theory to ground the theory of reference. Now all of a sudden we have turned the argument on its head: if this is the way reference relates to perception, perception must be providing a far richer range of entities than had previously been suspected. It is now a challenge for perceptual theory to describe how the perceptual systems accomplish this. In other words, examples like (11)–(12) open the door for fruitful cooperation between linguistic semantics and research in perception.

Glancing back at an issue raised in Chapters 5, 8, and 9, consider the degree

[18] Given that the use of *six Mahler's Seconds* is confined to a small community, I suspect it is shorthand for *six performances of Mahler's Second*, i.e. it is a conventionalized *coercion*, parallel to the more widespread *three coffees* for *three cups/portions of coffee*. See section 12.2.

to which syntactic distinctions mirror semantic distinctions. It turns out that the correspondence, such as it is, pertains above all to ontological category features. For instance, sentences map into situation- and event-concepts; PPs map primarily into place-, direction-, and path-concepts; APs map into property-concepts; and NPs map into just about anything. Going in the other direction, object-concepts map invariably into NPs, place-concepts map into PPs or NPs, and so on. This mapping is evident in the syntactic categories of the deictic forms in (12). In other words, ontological category features, unlike most descriptive features of concepts, *are* "visible" to the syntax–semantics interface. This means that they play an important role in grammar as well as in perception.

The repertoire of ontological types seems to me another good candidate for a very skeletal unlearned element of cognition. Again, because it is central not just to language but to perception and action, we need not call it part of Universal Grammar, the human specialization for learning language. But we do have to call it part of the innate human cognitive endowment.

One last comment before moving on. We have developed here a repertoire of major feature types: indexical features, valuation features, modality, descriptive features, and now ontological category as well. This classification, it seems to me, reveals a major sociological split among approaches to semantics. Formal semantics, growing out of the logical tradition, has concentrated on indexical features and valuations (which, as we will see in section 12.4, include negation, quantification, possibility, and attributions of mental states). By contrast, most traditions of lexical semantics, including cognitive grammar, have been primarily concerned with the descriptive features. A full semantic theory, of course, has to account for both.

10.9 Proper names, kinds, and abstract objects

So far we have been primarily considering referential uses of deictic expressions. Using the machinery developed for these cases, we can begin to construct the basics of other important elements of natural language semantics.

10.9.1 Proper names

It is standard to say that a proper name denotes an individual. In the present approach this translates into saying that a proper name has an indexical in its associated concept.

We can now attribute the distinction between names of real people like *Bertrand Russell* and names of fictional people like *Sherlock Holmes* to the valuation attached to the name: *Russell* has the valuation *external* and *Holmes* has the valuation *imaginary*.

We may also note that there are proper names for entities of other ontological categories. For example, *World War II* names an event, *Wyoming* a place, *Symphonie Fantastique* a sound, *1946* a time period, and *General Motors* an organization (whatever sort of ontological category *that* is).[19]

10.9.2 Kinds

All the concepts we have discussed so far have been individuals. Now consider the possibility of a cognitive structure that has descriptive features (including modality and ontological category), but lacks an indexical and a valuation. I suggest that this is just the structure we need for *kinds* or *types*.

To see the virtues of this approach, consider first the alternative possibility that the difference between an individual (or token) and a kind is that the latter is less specific in its descriptive features. The difficulty here is there are kinds whose descriptive characteristics are very sharp indeed—for instance manufactured objects such as 1998 Lincoln pennies, copies of the Boston *Globe* for April 4, 2000, and size 6-32 half-inch round-head brass bolts. What a kind really lacks is the possibility of pointing to it: one can only point to *instances* of it. Omitting the indexical feature from the concept would have exactly this effect.

This yields a nice formal consequence: one can form an instance from a kind simply by adding an indexical feature. This is paralleled in linguistic expression. Suppose that (as most semanticists think) a common noun such as *penny* denotes a kind, and that, as we have seen above, a demonstrative such as *this* or *that* expresses an indexical feature. Then it is no surprise that the full NP *that penny* denotes an individual that instantiates the type expressed by *penny*—semantically it is the unification of the descriptive features of *penny* and the indexical feature of *that*.[20] Conversely, given an individual, one can form a kind of which it is an instance, simply by deleting the indexical. This does not fall into

[19] Something should be said about the "causal theory of reference" for proper names (Kripke 1972). It is obvious that named entities normally get their names by virtue of some agentive event of naming; we assume a named individual has such an event in his/her/its history. I don't see anything especially profound here. I am inclined to think that having a name is conceptualized as a property of an object—a descriptive feature—not unlike having a size or shape. And of course houses get their size and shape by virtue of some agentive act of building; we assume a house has such an event in its history. But perhaps I am being too flip.

I cannot resist a personal anecdote that reveals something about children's implicit causal theory of reference (though I'm not quite sure what): each of my daughters, at about the age of 5, asked me, "If there weren't people when the dinosaurs were around, how do we know what they were called?"

[20] Quantifiers like *every* combine with common nouns in a different way, of course, since they do not express (simple) indexical features. See section 12.4.

linguistic expression quite so readily, but an apt expression is *things like that [pointing]*.

This relation of kinds and their instances also leads to a simple formal description of type-learning from a presented example: the learner deletes the indexical feature from the associated percept (in addition perhaps manipulating or thinning out the descriptive features). Such an approach parallels the approach to linguistic rule-learning advocated in Chapter 6: constructing a rule from an instance involves replacing one or more constants in a structural schema by variables that are now free to be used in a wide range of instantiations.

Language is not always entirely strict in marking the type–token distinction. For instance, *the same* can be used to express either token or type identity, with the choice being only a matter of pragmatics:

(14) a. John wore the same hat he always wears. [probably token-identity]
b. John ate the same sandwich he always eats. [probably type-identity!]

We would like this minimal difference to correspond to a minimal difference in conceptual structure. In the present analysis, it is a simple matter of letting *the same* be indifferent to whether an indexical feature is present or not. Similarly, an NP can denote an individual or a kind depending on its context in the sentence. In particular, a *predicate NP* is one that is not associated with an indexical feature.

(15) a. A professor walked in. [*a professor* = individual]
b. John has become a professor. [*a professor* = kind][21]

(15b) does not name two different individuals, John and a professor; rather it says that the descriptive features of professors have come to be associated with the individual John.

Most kinds come without valuations. One can perceive a real bug or experience an imaged bug, and they are both bugs. But some kinds do carry inherent valuations. Unicorns, for instance, are specified as mythical beasts, so the kind *unicorn* must carry the marked valuation *imaginary*.

And kinds come in various ontological categories. In addition to object types such as those expressed by *cocker spaniel* and *6-32 half-inch brass bolt*, there are event types, such as those expressed by *my eating breakfast* (on an unspecified occasion) and *its having rained yesterday* (in an unspecified place). Similarly, in

[21] Of course some other languages do mark this difference, as in French *Un professeur est arrivé* ('a professor has arrived') vs. *Jean est devenu (*un) professeur* ('John has become (*a) professor').

addition to specific locations such as those expressed by *in that cup over there*, there are location types such as those expressed by *in people's ears*.

On the other hand, some of the other ontological categories do not so clearly display a type–token distinction. For instance, I find it odd to think of the space between my hands and the length of the fish in (12e) as different instances of the general type expressed by, say, *nineteen inches*. I also find it hard to imagine a type–token ambiguity like that in (14) for sentences about distances, such as *John stayed the same distance from me he always stays*. This suggests that perhaps this ontological category formally lacks the possibility of an indexical feature. For a different case, the last section mentioned the ambivalent behavior of sounds, twinges of pain, sunrises, and Tuesdays with respect to the type–token distinction. I leave the question open as to how these phenomena are to be treated formally.

I want briefly to compare this treatment of kinds with the basic treatment in predicate logic and its derivatives, including many versions of formal semantics. There an individual is formalized as (or named by) a constant and a kind is formalized as a predicate with a variable. An individual that belongs to a particular kind is formalized in terms of an existential quantification over the predicate in question. So, for instance, the kind expressed by *professor* is notated as a one-place predicate *Px* ('is a professor'), and we get logical expressions like (16).

(16) a. John is a professor = P(J)
b. A professor walked in =
i. $\exists x(Px \,\&\, Wx)$
'there exists an x such that x is a professor and x walked in'
or
ii. $\exists x_{Px}(Wx)$
'there exists an x, x being a professor, such that x walked in'

An advantage of this notation is that it is hallowed by tradition and all practitioners know how to read it. I see three drawbacks. First, it tends to muddy the distinction among ontological categories: the object type expressed by the noun *professor* is always rendered logically by 'x is a professor', which in the present system is a type of state-of-affairs. Second, this notation makes it difficult to see the similarity between the two uses of *same* in (14). Third, once we recognize that language makes reference not just to objects but to events and places, it is necessary to clutter logical expressions with existential quantifiers. For instance, (17a) expresses an event containing two objects and a place, so its logical treatment has to be something like (17b).

(17) a. A boy sits on a chair.
b. $\exists e \exists x_{BOYx} \exists p \exists y_{CHAIRy}$ (e = SIT(x,p) & p = ONy)

The event-variable *e* is now commonly accepted since being championed by Davidson (1967); the place-variable *p* is as far as I know a neologism.

These objections to standard notation are not fatal by any means—it can, with appropriate tinkering, code the same set of distinctions. In any event I see no reason not to experiment with alternative notations in an effort to find something more perspicuous and more easily adaptable to a wide range of linguistic phenomena. The next two chapters will work out some alternative possibilities.

In denotational theories of meaning such as Kripke (1972), Montague (1973), and Putnam (1975), kinds are taken not to be f-mental structures consisting of descriptive features, but rather sets, such as the set of individuals (in all possible worlds) who are professors. This cannot be maintained in a mentalistic theory, as one does not carry the members of such sets around in one's head. Rather, at best one can carry a list of some prominent instances, plus a schema that allows one to generate and identify members of the set (i.e. instances of the kind)—if you like, a functional or intensional delineation of the set. But that is exactly what the descriptive features are for. The extension of the set of professors, the collection of its members, plays no role in the f-mind.

A justification for positing the set is the presumption that the reference of a phrase should be built up compositionally from the reference of its words: the set provides a reference for the common noun *professor* that can undergo Boolean combination with other sets. In the present approach, by contrast, *professor* simply has a sense without a reference. The *sense* of a phrase is built up compositionally, and if the composition happens to provide an indexical feature, then we get a referring phrase. I don't see any particular harm in this stance, as long as the indubitably referential phrases are taken care of.[22]

10.9.3 *Abstract objects*

Up to now we have been talking mostly about perceivable entities. What makes them perceivable is that they have features that connect to the perceptual interfaces in Fig. 10.5. Now suppose that conceptual structure contains other features that do not pertain to a perceptual modality, but which connect instead to the inferential system in Fig. 10.5. Such features would provide "bridges" to other concepts but no direct connection to perception. That is, they are used in reasoning rather than in identification. Let's call these "inferential features," by contrast with perceptual features.

What would be candidates for inferential features? Consider an object's

[22] Note that although *professor* is on its own non-referential, the phrases *the word 'professor'* and *the kind 'professor'* are referential, referring to a word and an abstract object (the meaning of the word) respectively.

value. This is certainly not perceivable, but it influences the way one reasons about the object, including one's goals and desires concerning the object. Value can be established by all sorts of means, mostly very indirect. Thus value seems like a prime candidate for a type of inferential feature. Another candidate is high-level taxonomic category. For instance, the classification of a particular object as an artifact predicts only that people typically use this object for something. But it does not predict what it looks like, how it works, or what one does with it. Rather, beyond the ontological feature 'physical object', which does have perceptual consequences, *artifact* is not on the whole a collection of perceptual features; rather it bridges to all sorts of inferences that can in turn potentially engage perceptual features.

Now suppose there were a concept that had *only* inferential features. This then would be an abstract concept such as those expressed by *the value of my watch* and *the meaning of this word*. From the point of view of language, abstract objects cannot be distinguished from concrete objects: all that is necessary for reference is an indexical feature. The interfaces to language cannot "see" whether a concept's descriptive features are perceptual or inferential. Thus as far as language is concerned, reference proceeds in the same fashion whether or not there is a conceptualized physical object out there in the world.

If abstract concepts have indexical features and descriptive features, in principle they should have ontological category features and valuations as well. An interesting issue for future research is to explore what ontological category features might pertain to abstract entities. As for valuations, compare *belief* and *opinion* with *value*, *meaning*, and *obligation*. The former two are understood as "subjective": they are objects in people's minds, hence *internal* relativized to a subject. By contrast, the latter three are normally understood as "objective parts of the world"—even though on reflection nothing has value without someone to value it, nothing has meaning without someone to understand it, and there are no obligations in the absence of individuals who understand the social system. That is, talking about someone's beliefs is like talking about their mental images or dreams (Jackendoff 1983: ch. 11); while talking about someone's obligations is more like talking about their car (Jackendoff 1999). In short, what causes something to be judged "real" is its having the valuation *external*—even when it is a patently abstract object![23]

It is a fascinating question why values, meanings, and obligations should be understood as "objective," and I am not in a position to answer it here.

[23] Thus an individual's having a "theory of mind" (in the sense of Premack and Woodruff 1978 and Wimmer and Perner 1983) could functionally depend on their being able to recognize the valuation *internal* as attributed to other individuals. The literature suggests that children do not develop this ability until about 5 years of age, and that most animals lack it almost entirely.

However, if indeed they are, it helps explain the intuitions behind the realist stance on language, in which words are "out there in the world," meanings are "out there in the world," and therefore our job is to explain how they connect to the rest of the world. Here, instead, we have asked why words and meanings "present themselves" to us as "out in the world," why "the world" presents itself to us the way it does, and how the cognitive connection between language and the experienced world is established. This is a more roundabout route, but in the end I hope a more fruitful one.

Returning for a moment to abstract concepts, it is worth asking whether animals could have them at all. Social concepts seem to me to have the right flavor. Consider the sophisticated way that monkeys operate with kinship relations and dominance hierarchies (Cheney and Seyfarth 1990). There is no space to go through the arguments here, but I am persuaded that these notions are in no sense directly visible "in the world," reducible to simple visual (or, say, olfactory) stimuli. Rather, they are abstract relationships that bridge a vast number of different inferences and behaviors. Such concepts are no longer "cheap tricks" that govern behavior; they are "expensive tricks" that coordinate a large number of disparate behaviors. There is no doubt that humans have many more abstract concepts than animals, and that language is instrumental in the exuberant proliferation of such concepts. But abstract concepts do not develop out of nowhere: the evolutionary precedent is there.

10.10 Satisfaction and truth

An important piece has been left out of the story so far. Linguistic expressions used in isolation cannot refer: they can only *purport* to refer. For example, suppose I utter *Hey, will you look at THAT!* to you over the telephone. You understand that I intend to refer to something; but you cannot establish the reference and therefore cannot establish the contextualized meaning of the utterance.

A referential expression *succeeds* in referring for the hearer if it is *satisfied* by something that can serve as its referent. In realist semantics, satisfaction is a relation between a linguistic expression and an entity in the world; in conceptualist semantics, the entity is in [the world as conceptualized by the language user]. It is this latter notion of satisfaction that must be worked out here. Lest this notion be thought too subjective and/or solipsistic, I wish to remind the reader of the commitments of the conceptualist program set out at the beginning of Chapter 9. We are interested in how people (including ourselves) understand the world, not in some notion of absolute truth.

To work out a conceptualist notion of satisfaction, we invoke one component

in Fig. 10.5 not yet discussed: the integration of concepts with one's f-knowledge base. Suppose I say to you: *I talked to David Wasserstein today*. In isolation, the proper name *David Wasserstein* purports to refer; you assume that I intended it actually to refer. If you know David Wasserstein (i.e. he is part of your knowledge base), you can establish the intended reference in your own construal of the world. If you do not know him, the proper name is unsatisfied *for you*.

More generally, four different situations can arise.

(18) a. The purported referent is present in your f-knowledge base or the readily available context, and it satisfies the referential expression.

b. The purported referent is not present in your f-knowledge base or the readily available context, and so you add a new character into your f-knowledge base to satisfy the referential expression. If I anticipate this situation for you, I will use a specific indefinite expression like *a friend of mine* or *this friend of mine, David Wasserstein* rather than a definite description or an unadorned proper name.

c. The purported referent is in conflict with your f-knowledge base, as in Russell's famous example *the present king of France*. Or it is in conflict with the readily available context, for example in a situation when I speak of *the apple on the table* and you see either no apples or two apples on the table. In this case you have to fall back on some repertoire of repair strategies (H. H. Clark 1996: ch. 9): guessing what I intend,[24] asking me for clarification, deciding to ignore me altogether, and so on.

d. The purported referent contains descriptive features that inherently conflict with each other, so that nothing can possibly satisfy the expression. In such a case, for instance *the square circle*, you find the expression anomalous, and must again fall back on repair strategies.

This is all clear with the reference of NPs. We next turn to the reference of sentences. The standard custom in the formal semantics tradition, going back to Frege (1892), is that the intended referent of a (declarative) sentence is a truth value. I must confess I have never understood the argument for this position (e.g. Chierchia and McConnell-Ginet 1990: ch. 2); to unravel it is beyond the scope of the present text (see however the discussion in Seuren 1998: 375ff.). I

[24] This would include the well-known situation described by Donnellan (1966), here adapted to the conceptualist approach: Joan speaks to Ida of *the man drinking a martini* and gestures toward Edgar. Ida, however, happens to know/believe that Edgar is drinking water, and hence cannot take him to satisfy Joan's description. Nevertheless, Ida will probably be charitable and understand Joan to intend the phrase to refer to Edgar. We do not have to worry about whether the phrase *really* refers, only about how language users treat it.

am going to explore instead the alternative position that the intended reference of a declarative sentence is a situation (an event or a state of affairs) (Jackendoff 1972). Traditionalists should not worry—truth values will get their due shortly.

The view that sentences refer to situations is motivated largely by linguistic parallels to referentiality in NPs—a kind of evidence not frequently cited in the literatures of philosophy of language and formal semantics. First notice the parallel in phrases that accompany deictic reference:

(19) a. Will you look at *that*! A blimp!
b. Will you look at *that*! The baby's walking!

(19a) draws the hearer's attention to an object in the environment; (19b) to an event—not to a truth value. Second, notice that discourse pronouns can corefer with sentences as well as with NPs.

(20) a. A blimp appeared in the sky. It was huge.
b. Suddenly the baby started walking. It astounded her parents.

The antecedent of *it* in (20b) is the whole preceding sentence, so it presumably has the same referent. But *it* certainly does not refer to a truth value: *it astounded her parents* does not assert that the parents are astounded by the truth of the proposition expressed by *the baby started walking*; they are astounded by the *event* of the baby walking. Hence this event must be the referent of the preceding sentence as well.

Next consider embedded *that*-clauses in a context where they alternate with NPs.

(21) a. The apple on the table astounded Max.
b. That the Red Sox won today astounded Max.

What astounded Max in (21b) was not a truth value; it was an event (a truly astounding one!). Parallel to the four possible situations in (18) for satisfaction of an NP's purported referent, there are four possible situations for the referent of *the Red Sox won today*.

(22) a. Your f-knowledge base includes the event of the Red Sox winning, and this satisfies the intended referent of the clause.
b. Your f-knowledge base does not include the event of the Red Sox winning, so you add this to the f-knowledge base as the referent of the clause.
c. Your f-knowledge base includes something in conflict with the purported event of the Red Sox winning (say your take on the world is that the Red Sox didn't play). Then your have to engage in some repair strategy.

d. The descriptive features of the purported event are inherently in conflict, so that there is no possible referent. In such a case, for instance *That the square is a circle astounded Max*, you judge the clause anomalous and again have to resort to repair.

There are of course other cases of *that*-clauses that do not come out this way, notably in so-called intensional contexts such as the complement of *believe*. However, NPs in this position are subject to the same distortions of referentiality:

(23) a. Max believes that there is a tooth fairy/the square is a circle.
b. Max believes in the tooth fairy/in square circles.

In both of these, the speaker makes no commitment to the existence of the tooth fairy or to square circles. In present terms, this context imposes a non-default valuation on the material contained in it, permitting reference to imaginary, nonexistent, and even anomalous entities. This valuation is indifferent as to whether the entities in question are objects or events. We take this up again in section 12.4.

So far, then, we have seen that it makes sense to regard a clause as referentially satisfied by a conceptualized situation. The judgment of a declarative sentence's truth value then follows from how it is referentially satisfied.

- In case (22a), where the intended referent of the sentence is present in the hearer's f-knowledge base (or can be deduced from it), the sentence is judged true.
- In case (22b), where there is no conflict with the hearer's f-knowledge base, the sentence is normally taken to be informative and presumed true.
- In case (22c), the sentence is judged false.
- In case (22d), it is judged analytically false.

Thus truth is defined in terms of reference and satisfaction, just as proposed by Sher (1996), following Tarski, in the quote at the outset of the chapter.

In short, the parallelism in the reference of NPs and sentences lies in the parallelism between conceptualized objects and conceptualized situations. The notion of satisfaction applies identically to both. However, sentences have an additional layer, in which they are characterized as true or false on the basis of how they are referentially satisfied.

Of course, truth and falsity pertain only to sentences that are asserted, that is, that have assertive illocutionary force—declarative main clauses. Many other sorts of illocutionary forces are possible, for example those in (24) (Austin 1962; Searle 1969; Bach and Harnish 1979; H. H. Clark 1996).

(24) a. Where is my hat? [Question]
b. What's the capital of New York? [Pedagogical or quiz-show question]
c. Would you please open the window? [Request]
d. Open that window, young man! [Order]
e. Shake well before using. [Instruction]
f. Suppose there were a bug on your leg. (What would you do?) [Supposition]
g. What a great theory this is! [Exclamation]
h. Have a great day! [Greeting]
i. Thanks for the lovely party. [Acknowledgement]
j. If only it would rain! [Invocation]
k. Let's take a vacation. [Suggestion?]
l. Let $x = 4y$. [Postulate/Axiom/Stipulation]
m. I promise to finish this book by June. [Promise]
n. One more step and I'll shoot. [Threat]
o. I name this baby Herman. [Dubbing]

There are also uses of simple declaratives that are not intended as assertions, that is, that are not intended to be referentially satisfied. Their truth value is not at issue in what the speaker intends to communicate.

(25) a. Once upon a time there was a little girl who lived in a forest. [Beginning of story]
b. There's a priest, a minister, and a rabbi in a boat . . . [Joke]
c. Joe can eat 3 candy bars in a minute. (How many can he eat in an hour?) [Pedagogical problem]
d. Your mother wears army boots! [Ritual insult (cf. Labov 1972)]
e. We all live in a yellow submarine. [Song lyrics]

For present purposes it is not too important to sort out which of these are distinct cases and what the full repertoire is. The only point, going back to Austin, is that making assertions that can be true or false is only one of many things we can do with language. Theorists' concentration on truth value, which seems to go all the way back to Plato, blinds us to the full vivid range of possibility.[25]

In the present approach, the problem of characterizing the conditions under which a sentence is judged true does not go away. It is just demoted from the

[25] People sometimes say that we need to deal with truth as an important part of the problem of meaning, then go on to ignore all the rest. And some (e.g. Katz 1977 for instance), recognizing the other possibilities, attempt to construct the meanings of sentences of all non-assertive illocutionary forces on the basis of truth values of corresponding assertions.

paramount problem of semantic theory to one among many problems. What seems more basic here is the conditions of satisfaction for referential constituents and how they interact with the f-knowledge base. In fact, a great deal of research in "truth-conditional" semantics can easily be reconstrued as addressing this issue. For instance, the question of whether sentence S_1 entails sentence S_2 has nothing to do with their truth values—they may describe thoroughly fictional or hypothetical situations. Rather, S_1 entails S_2 if adding the situation referred to by S_1 to an otherwise blank f-knowledge base enables the situation referred to by S_2 to be satisfied. The factors involved in such satisfaction, and the form of the rules of inference, may remain essentially unchanged from a truth-conditional account.

Above all, the conceptualist approach shifts the focus of semantics from the question "What makes sentences true?" to what I take to be the more ecologically sound question, "How do we humans understand language?"—"ecologically sound" in the sense that it permits us to integrate semantics with the other human sciences. I take this to be a positive step.

10.11 Objectivity, error, and the role of the community

We're not yet quite done escaping the accusation of solipsism. A good entry into the remaining difficulty is through the problem of error. The problem is usually stated something like this:

(26) a. Joe says "Look at that duck!"
b. But in actuality he points to a platypus.
c. How does his phrase *that duck* manage to refer?
d. Does *duck* still refer to ducks in (26a)?

Within common-sense realist theories of meaning, in which words refer to the world, the problem lies in (26c, d): how the phrase could refer to something it is not supposed to, or how the phrase could get its meaning despite being on occasion misused.[26] Within a conceptualist theory this is not a problem: in Joe's conceptualization of the world, that thing "out there" is a duck, so for Joe the phrase succeeds in referring.

The problem for the conceptualist theory lies rather in sentence (26b): *But in*

[26] Famous cases in the literature that have the form of the argument in (26) include Putnam's (1975) "cats turn out to be robots radio-controlled from Mars" scenario; Putnam's (1975) "Twin-Earth" scenario; Donnellan's (1966) "'the man with the martini' is really drinking water" scenario; and Kripke's (1979) scenario of Pierre who thinks London is ugly and "Londres est jolie." My case of the two Blooms (section 10.7) could also be couched in this form.

actuality he points to a platypus. By pushing the world into the mind, we no longer have a standpoint in objective "actuality," independent of the minds of Joe and the other people observing the situation. All we can say is something like "But in the conceptualized worlds of other observers, he points to a platypus." What is to stop them from all having different conceptualizations? How can we speak any more of Joe's being mistaken? And why should we care? Are we therefore doomed to a radical relativism about knowledge?

A way out comes from recognizing the sequence (26a, b) not as a disembodied piece of language "out there in the world," but as originating from a speaker who is communicating his or her conceptualization of the world. This conceptualization includes the judgment that Joe said something with intent to refer to some entity, and that this entity was a platypus. And now we can ask how the speaker came to this judgment.

The judgment has to be based on the match between two concepts: the descriptive characteristics of the entity in question and the descriptive characteristics associated with the word *platypus*. How did the speaker acquire these concepts? The characteristics of the entity Joe names may have come through observation of the reported scene, or they may have come through hearing someone else's report of such an observation. The characteristics of the kind *platypus* have to come through the speaker's learning the word. This in turn relies on a history of observation of others' usage (including instruction and demonstration).

I belabor this point because it poses the same difficulty for conceptualism as (26) itself. How can anyone rely on someone else's report of an observation, and how can anyone count on others' usage of language? Turned a different way: if conceptualization is essentially personal, how can we communicate?

This is now getting quite speculative, but I am inclined to think that human beings have a need to "tune" their conceptualizations to those of others—to have a common understanding of the world. Since we cannot read minds, the only way we can assess our attunement is by judging whether the behavior of others—including what they say—makes sense. To the extent that all members of a community are effective in acting on this drive, there will be a tendency for conceptual convergence, including in the use of language.

A learner faced with a community that has already converged will see the problem of tuning as one of adapting oneself to the community's conceptualization: "This is how things are done." On the other hand, if the community has not converged, the learner will be faced with some difficulty sorting out what to do. This is as true of cultural and technological learning as of language learning.

Of course there are many reasons people's conceptualizations do not converge. People may have access to different information, due to their different

experiences. Conflicts may arise because people have different goals and needs, which affect what they find important in any given situation. Conceptualization based on one's own observation or imagination may conflict with the "received" conceptualization with which one desires to tune oneself. The worst is probably when two groups with different communal conceptualizations come into contact, leading to massive breakdown of the sense of attunement.

There are various ways to seek attunement. One way is reliance on an "expert," as in Putnam's (1975) "division of linguistic labor." This includes of course how children learn their culture. When the information is social information, another way is gossip. And apparently the members of a community can dynamically "co-tune" each other, for example when the children of a pidgin-speaking community together converge on a creole that did not exist before.

An individual faced with a perceived lack of attunement has certain choices. One is to accede to the perceived consensus—"accepting the truth." Another is to attempt to tune others to oneself, "convincing them they are wrong," through the exertion of social coercion, whether explicit ("You're wrong. How can you be so stupid?") or implicit ("That approach isn't interesting."). A third way requires a brief digression.

Searle (1995) introduces the idea of a conceptualized *joint intention*. Typically we think of an intention as a cognitive structure that strictly concerns only the performance of one's own action. You can't intend for someone else to perform an action, except indirectly: by intending to do something yourself that brings it about that the other person performs that action (Jackendoff 1995). However, in a joint intention, the intended action is to be carried out by oneself plus others; one's own intended actions fulfill a role in the joint action. For instance, each member of a team has a joint intention to fulfill the team's goals, along with an individual intention to play a role within the team's joint action. Searle makes the case that any sort of deliberately cooperative behavior requires a conceptualized joint intention on the part of the participants, and that the operation of most social institutions relies on an understanding of joint intentions.

Herbert Clark (1996) argues persuasively that (in these terms) a speech act requires not just an intention on the part of the speaker to convey information, but a joint intention on the part of the speaker and hearer towards the conveyance of information. Most of his book is about the delicate negotiations and signals that go on during speech in order for the participants to keep in synchrony and to verify that the joint intention is being fulfilled. It is clear that he means this coordination as a special case of a more general phenomenon of cooperative behavior: he constantly compares conversation to tasks like moving furniture and playing chamber music. In fact it is hard not to see joint inten-

tion in the collaborative actions of highly social mammals such as wolves and especially chimpanzees.

Establishing and maintaining a joint intention is yet another way of tuning conceptualizations. It differs from the other ways of tuning conceptualizations—either accepting or fighting the status quo—in that it is not one-sided: the participants are cooperating in establishing a tuning and actively maintaining it.

However, such cooperation is subject to defection. An all-too-common example occurs when one participant takes the common enterprise to be one of cooperation in solving a problem (i.e. presumes tuning), and the other takes it to be one of social coercion ("the truth is what *I* say it is"). The first participant, on discovering the defection, justifiably accuses the second of a swindle or worse.

This brief discussion of "tuning" may evoke in readers a rich range of associations, from education to politics to mental illness to multiculturalism to postmodernism. This is not the place to go on; I encourage others to explore the connections.

For present purposes, the point is that "tuning" seems to me the last piece we need to close the loop in creating a conceptualist account of reference and truth. Not only is our conceptualized world our own reality, we constantly check whether it converges with everyone else's. To the degree that we sense that it converges, we take the common view as flowing from the "objective character of the world." On the other hand, to the degree that we sense conflict, we are forced to acknowledge subjectivity, and the sense of what is "objective" becomes less stable.

This is not to say that objectivity and truth reduce to consensus (a view urged by Rorty 1979, for instance). Rather, objectivity is understood as an ideal that we aspire to achieve. Consensus or "co-tuning" is one important factor that the f-mind weighs (and often *we* weigh consciously) in judging objectivity and truth. But there are others. The rebel trusts his or her own perception and inference and holds out against the consensus ("The emperor has no clothes!"). Achieving the proper balance is often difficult, and the line between courage and madness can be a tough call.

Here is where I think we are, after pursuing a rigorously mentalist approach to meaning. We have not ended up with the rock-solid rigid notion of truth that the realists want. Rather, I think we have begun to envision something that has the promise of explaining our human sense of truth and reference—and of explaining why philosophical and commonsensical disputes about truth and objectivity so often proceed the way they do. I have no illusions that this work is over, but it strikes me as a path well worth exploring.