

**The Co-Ascription of Ordered Lexical Pairs:
a Cognitive-Science-Based Semantic Theory of Meaning and Reference.
Part 1**

Tom Johnston
tm1972.com
tmj44x@att.net
August, 2018

Table of Contents

The Co-Ascription of Ordered Lexical Pairs: a Cognitive-Science-Based Semantic Theory of Meaning and Reference. Part 1	1
Abstract.....	2
Introduction.....	2
A Methodological Note.....	4
Section 1. Linguistic Components.....	11
Concepts.....	11
Statements.....	13
Statement Tokens and Statement Awareness Events.....	18
Section 2. Meaning and Reference.....	20
Meaning: Concept-to-Concept Links.....	20
Meaning Sets.....	25
Reference: Concept-to-Object Links.....	29
Reference Sets.....	35
Indeterminate Reference: Association Sets.....	38
Meaning and Reference: Recursion.....	40
Meaning and Reference: Unification.....	42
Section 3. Measuring Meaning, Reference, and Semantics.....	45
Lexical Pairs and Canonical Forms.....	45
Measuring Meaning: Concrete Concepts.....	49
Measuring Meaning: Abstract Concepts.....	55
Measuring Reference: Concrete Concepts.....	56
Abstract Concepts as Concrete.....	62
Measuring Semantics.....	65
Conclusion.....	67
Bibliography.....	70

Abstract

Lexical semantics has a problem. As Alessandro Lenci put it, the problem is that it cannot distinguish semantic from non-semantic relationships within its data. (2008, 2014). The data it relies on are patterns of co-occurrence of lexemes within linguistic corpora. But patterns of co-occurrence can reflect either our knowledge of what the world is like or our knowledge of what words mean -- matters of fact or matters of meaning.

In this essay, I develop a semantic theory which draws this distinction in a way which makes it discernible in lexical semantics and cognitive science research. In doing so, this theory unifies truth-functional and structuralist approaches to semantics, and provides an integrated explanation of meaning and reference.

I base this semantic theory on linguistic dispositions to pair words with words, and to pair words with things, both based on learned patterns of association. These dispositions manifest themselves in verbal behavior, and (in Part 2) I propose a new neurophysical model to account for these dispositions which underlie both the intensional and extensional patterns in that behavior.

Introduction

Because of its origins in Saussurean structuralism, lexical semantics in general, and distributional semantics in particular, looks for meaning in relationships among words. Because of its origins in the development of formal logics, and the technique of giving interpretations to the axioms and theorems of formal systems by means of models which map elements of the systems onto referents in the world, analytic philosophy of language in general, and truth-functional semantics in particular, looks for meaning in the relationships between language and what language is about.

Figure 1 shows semantics as consisting of two components: meaning and reference. Two of these terms -- "semantics" and "meaning" -- are often used interchangeably, but I don't use them that way. Instead, I use them like this. Semantics is the set of constraints on words which, when combined in statements, convey information. One subset of those constraints is meaning. This is what post-structuralist disciplines such as lexical

semantics concentrate on. The second subset is reference. This is what analytic philosophy of language concentrates on.




Figure 1. Semantics as Meaning and Reference.

The roots of reference (to borrow a phrase from Quine) are the constraints which shape the ascription of a concept to a range of sensory inputs. The roots of meaning are the constraints which the ascription of one concept places on the ascription of another concept, forbidding, requiring or permitting the second ascription. The question is what those constraints are.

Both constraints work in the context of statements. Thus, statements are both meaningful (in the structuralist sense) and referential (in the truth-functional sense).

Research in distributional semantics into lexical-pair co-occurrences has failed to account for this distinction. Research in the cognitive neuroscience of language often fails to make this distinction at all in the design of experiments or in the interpretation of data.

Although meaning and reference are distinct semantic forces, they mutually constrain one another. In doing so, they create the information content of statements. Another important fact about meaning and reference is that they have a common behavioral and neurological foundation. Both meaning and reference are the result of dispositions to behave linguistically in patterned ways. With both, those dispositions are the result of ever-evolving associations of neurally-encoded patterns in human brains whose resulting patterns of linguistic behavior converge, across intersecting groups of speakers, to an

extent sufficient to enable us to exchange information.

So there is a scientifically-grounded semantic theory to be had. TM, as I call it, is such a theory.

A Methodological Note

No scientific theory exists in isolation from a set of beliefs which provide a context within which it is related to other theories, and without which it cannot even be expressed. Every formal axiomatic system of predicate logic, for example, needs a model which interprets its basic axioms, their derivative theorems, and the predicates with which those axioms and theorems are expressed. Without that model, the system is pure mathematics, about nothing at all except the patterns which that mathematics reveals.

This set of beliefs thus constitutes a model, as it were, one which is involved in the contextualization and comprehension of more specialized models. The most fundamental and all-embracing of these models is *background theory* -- commonsense beliefs shorn of inconsistency and of the burden of idiosyncratic beliefs which the uninformed and unthoughtful are only able to retain because they aren't aware of the rich web of other beliefs with which their outre beliefs are inconsistent (flat-earth believers, for example), and because they have no interest in or skill with formal or informal processes of truth-preserving deductive inferences or probability-supporting inductive inferences. This cleaned-up version of commonsense is the *background theory* we all more or less share.

I represent this model of all models, this background for everything we think about, as the top-most box in Figure 2. Specific theories, both in science and in other fields of particular study, emerge out of this background theory, and would be unintelligible without it. This theory contains beliefs that are expressed in statements that we all, more

or less, take to be true.

In particular, the collection of beliefs we have about our mental lives has been called our "folk theory of mind". Rather than considering these beliefs as consisting in large part of mistaken superstitions (which is the eliminativist position of Patricia and Paul Churchland), I consider these beliefs as setting the problems which scientific theories attempt to solve.

The neuroscientist Stanislas Dehaene said that "The correct perspective is to think of subjective reports as raw data". (2014; p.12). Such subjective reports are made in the language of our background theory, and in particular of its folk theory of mind component. Like Dehaene, I think that these reports constitute data, the material which must be accounted for.

Of course, all data is already interpreted, and so is subject to redescription which renders it consistent with evolving theory, or sometimes to rejection. But for the most part, well-established background theory beliefs about mind and about language are what our theories about mind and language have to explain. This is not because such well-established beliefs are somehow true from some God's eye point of view. That could never be established. It is because unless or until we reformulate how we express our basic beliefs, that language is how we report on what we encounter. It is our "observation language", the "go to" language we use for expressing what it is we want to further explain with our scientific theories. If there were a God's eye point of view from which fundamental beliefs were quite different, that wouldn't matter.

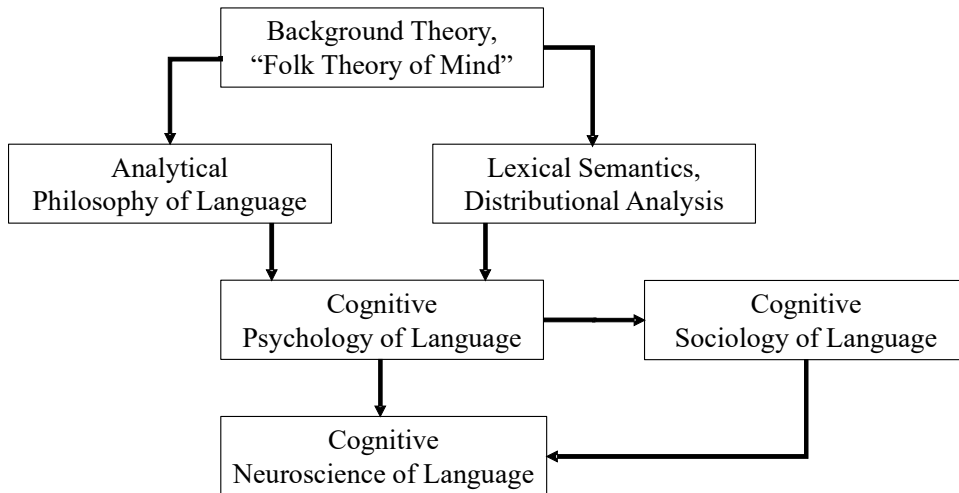


Figure 2. Theoretical Perspectives on Language.

So as I develop theoretical accounts of language in each of the conceptual frameworks indicated by the five other boxes in Figure 2, the reader should remain aware of how extensively our background beliefs and observation language permeate these more theoretically special descriptions. As for the five theoretical perspectives themselves, their dependency relationships should be viewed as analogous to well-known dependency relationships in more fundamental physical science: that cell biology, for example, must be a consistent application of molecular biology, and it of organic chemistry, and it of physical chemistry, and it of the physics of the atom.

Because my topic is semantics and its meaning and reference components, the next two boxes down in Figure 2 represent the two very different approaches to semantics which I described above. The philosophical perspective has always been predominantly about reference, and in particular about the relationship between true statements and what in the world it is that they refer to and that makes those statements true.

Certainly, the concepts used in statements have something to do with the truth of those statements, and so the focus on statements and their truth cannot ignore concepts. But by

and large, some correlative to a correspondence theory of truth -- such as Kripke's causal theory of reference -- is presented, at which point the focus of attention shifts back to statements and truth, and away from concepts and meaning.

On the other hand, the lexical semantics perspective has always focused on the mutuality of the semantic influences which lexemes exert on one another. A structuralist web of lexemes is said to determine the meaning of each item within that web. Very little attention is paid to the relationships between statements and what in the world they are about, and thus very little attention is paid to reference.¹

If language is considered as an abstract complex object -- which is a perspective which follows from taking a synchronic approach to the study of language -- then I see no way to move beyond the interminable discussions and perspectives and debates about semantics which still characterize both current philosophical and linguistic approaches. Appeals to linguistic behavior, for example to pairs of sentences which differ in very minor ways but for which one sentence is judged to be acceptable and the other aberrant, are made as and when needed to support some synchronic perspective contention. But whether or not we call recent work in these two fields "progress" over what went before, this work has not gotten us to solid ground. Consider Chomsky's "progress" in the study of syntax (and, eventually, semantics): phrase-structure grammar, transformational-generative grammar, principles and parameters, x-bar theory, and the minimalist program, a series in which Chomsky explicitly repudiated much of his earlier work including such important methodological principles as the competence/performance distinction, and the criterion of simplicity for evaluating theoretical proposals, which he later replaced with the criterion of maximal constraint. (Harris, 1993; pp. 96-100, 176-187.)

Instead, I begin with the basic fact that language, as a synchronic object, is an abstraction

¹ For an insightful account of these two approaches to semantics, see the *Introduction*, by Santambroglio and Viola, to (Eco, Santambroglio and Viola, 1988).

from patterns manifested in public acts of linguistic behavior, acts in which sentences are produced or comprehended. Each such act occurs on a particular occasion, and is the act of a particular person. This way of studying language, and especially semantics and its meaning and reference components, is the province of the cognitive psychology of language, the next box down in Figure 2.²

The box to the right of the cognitive psychology box, in Figure 2, represents the conceptual perspective of what I call the cognitive sociology of language. For this perspective, the central issue is how individual patterns of linguistic behavior, which certainly vary from person to person and, orthogonally, over time, nonetheless become similar enough that when we speak, we aren't usually unintelligible to whomever we are speaking to. Somehow, my way of using language influences how others within my language communities use the same language, and vice versa for all of us.

A successful explanation of this fact, I believe, will explain what it is we are talking about with our reified vocabulary of language as an object to be studied, and of sentences and statements as objects of investigation, rather than what it is we are talking about when we discuss tokens of those abstractions, tokens which are produced and comprehended, each time, by one person on one occasion. It will demonstrate that when we take that synchronic perspective, and talk about such abstract objects as language itself, sentences, concepts, lexical categories, parts of speech and so on, we are talking about generalizations of patterns exhibited in large sets of utterances, as derived by formal and informal lexicographic and grammatical generalizations and, in the former case, as recorded in etymological dictionaries and in the periodic revisions of standard

² The "cognitive" here is used to indicate that most of the data I will be concerned with has to do with processes involved in the production and comprehension of language that we are or could be consciously aware of, e.g. with consciously speaking, writing, listening to and reading sentences. To illustrate the distinction, the construction of phonemes out of phones, or syllables out of phonemes, is part of the psychology of language, but is not part of the cognitive psychology of language, as I use the term in this essay. Note that this is a narrower sense of "cognitive" than the one in general use in the cognitive sciences.

dictionaries of natural languages.

Linguistic behavior is the behavior of human beings and, as such, is the output of states and processes of human brains. This perspective takes us to the bottom-most box in Figure 2. Contemporary neuroscience has advanced to the point where the cognitive neuroscience of language is now an established and flourishing subfield. But, in my opinion, these neuroscientists cannot see the forest of semantics clearly because they cannot clearly see the trees which make up that forest. They often use the term "semantics" in a way which ignores the distinction between meaning and reference as in, for example, the design of experiments which include both word-to-word matching and also word-to-picture matching. To the degree that their statistical compilations of the data from such experiments, and their theoretical hypotheses explanatory of those findings, do not clearly separate word-to-word conclusions from word-to-picture conclusions, let alone address the issue of how these two distinct background theory, philosophical, linguistic and psychological phenomena are differently implemented in the human brain and how they relate, neurally, in a way which explains how they relate at these other levels, then their research will continue to be hampered by the interference of data about meaning with data about reference (and vice versa), thus severely hampering the ongoing attempt to create a perspective which tells us about the structure and evolutionary behavior of the forest in which those different kinds of trees are found. On this point, I conclude (in Part 2 of this essay) by offering a set of extensions to and modifications of such well-known neurophysiological semantic theories as the Wernicke-Lichtheim-Geschwind Model, the Dual-Stream Model (Hickok & Poeppel, 2007; Hickok, 2009), the Dual-Coding Model (Paivio, 2007), the Lemma Model (Levelt, 2001), and the Hub-and-Spoke Model (Hoffman, et. al., 2011), all of which are discussed in (Kemmerer, 2014).

So these are the lines along which I will develop TM, in the process accounting for the well-known "encyclopedia knowledge vs. dictionary knowledge" distinction which

expresses the problem that lexical semantics has in developing a semantic theory. Finally, TM reconciles the different approaches to semantics of philosophy and linguistics, the analytic truth-functional approach and the structuralist lexical network approach.

Section 1. Linguistic Components

Concepts

Concepts are sub-sentential units of meaning.³ For Saussure, and post-structuralists, meaning is to be found in the constraints which concepts place on one another. But these constraints manifest themselves in the context of statements made by specific people on specific occasions; and in those contexts, they refer to the things those statements are about.

I need to make the notion of a concept more precise, and I begin with the background theory notion of a *category*. A category is a group of things that are alike in some way. This has both similarities and dissimilarities with the mathematical notion of a set.

To begin with, any non-random set is a collection of things for which there is a set membership criterion. So something is a member of such a set if and only if it is picked out by the set's membership criterion. The same is true for categories, but with two important differences. First of all, the membership criterion for a mathematical set must be precise enough to determine, for every candidate member, i.e. a member of the universe of discourse for that set, whether it is or is not a member of that set. But categories don't have that requirement. For example, at what point on a temperature gradient does a warm object become a hot one? Do the categories [Motor] and [Engine] have the same members?

The second difference between sets and categories is that mathematical objects do not change over time. Add or remove a member to a mathematical set, and the result is a new set. But new members can be added to a category, as and when discovered, and existing

³ Here I use "meaning" in the sense of meaningfulness, of having semantic content. "Significance" is a near synonym of "meaning" used this way.

members can change in ways which make them no longer members of that category, and the category can remain the same category in spite of these changes in membership.

So categories don't have precise and time-invariant set membership criteria. They have, instead, *fuzzy category membership criteria* which may change over time in anticipation of or in response to acts of including or excluding specific things or features in the acts of referential use of those categories.

I can now give a preliminary definition of a concept. A *concept* is a category which is a sub-sentential unit of meaning. So whether or not a concept applies to some object or other may not be clear-cut. In addition, a concept may remain the same concept in spite of things ceasing to be or coming to be instances of that concept. The concept of a horse does not change as some horses die and new horses are born.

Concepts are physically expressed as words or phrases that are spoken, or written, or represented in some form of sign language. I use "concept" synonymously with "lexeme", as understood as follows. A string of letters or sounds which may correspond to a word (or morpheme, or multi-word expression) is not a lexeme if its semantics is not involved in its production or comprehension. A lexeme is the sound+meaning (or orthography+meaning) of a sub-sentential expression in language. And that is what a concept is.

An *object* (more informally, a *thing*) is not a concept. It is a specific thing, whether physical, mental, or abstract. If it is a mental or abstract thing, however, the physical word itself is the only physical representation associated with it.

An important distinction between things and concepts is that concepts have instances. These instances can be either other concepts or things. But things don't have instances;

they *are* instances. As cognitively accessible, they are instances of concepts. In the language of classical philosophy, concepts are *universals*, and things are *particulars*.⁴

I distinguish two roles that concepts play. In one role, they are *kinds*, indicating the kind of thing something is, or a more general category under which another concept falls. That thing over there is a tree; the thing next to it is a rake. In this role, the concepts [Tree] and [Rake] are instantiated as particulars (or objects, or things). In a second role, concepts are *features*, indicating something about the thing picked out. That tree is beautiful; that rake is broken. [Beautiful] and [Broken] are features. In this role, they are instantiated as properties of things. Features may also be instantiated as relationships things have to one another.

In their role as kinds, concepts help to pick things out. In this role, they are accompanied by indexicals, linguistic or gestural devices which designate an object of the indicated kind. In their role as features, concepts say something about what has been picked out.

Statements

In language, information is expressed in *statements*. Some statements are combinations of other statements. For example, "Roses are red and violets are blue" is a compound statement, a conjunction of two simpler statements. Co-ordinate, superordinate and subordinate clauses, and adjectival or adverbial words and phrases, may also add information to a statement, which then becomes a complex statement. But a *basic statement* is expressed in a simple sentence, not a compound or a complex one. A basic statement picks out one thing and says one thing about it.

A basic statement derived from simplifying a compound and/or complex statement is

⁴ See (Loux, "The Problem of Universals", and "Particulars and Their Individuation". In Loux (ed), 1970). See also (MacLeod and Rubenstein, 2018).

often not information-content-equivalent to the statement it is derived from. For example, the basic statement list {"Roses are red", "Violets are blue"} may be a decomposition of either a conjunction or a disjunction; and those two compound statements do not express the same information. The issue here is the semantics supplied by syntax; and by focusing on basic statements, I exclude that issue from this essay.

By virtue of picking something out and saying something about it, a statement is an indicative sentence. But more than that, it is a *declarative sentence*, one in which the speaker expresses what she believes to be true. This belief is a *propositional attitude* the speaker has to what she has said. Gricean implicatures then lead listeners to understand the statement as a declaration of what she believes to be true.⁵

In making a basic statement, a speaker ascribes a kind to an object and a feature to a property of that object. In doing both, the speaker co-ascribes a pair of concepts to a *referent*. These individual acts of ascription depend on and express the *reference* of each of the two concepts. This act of *co-ascription* depends on and expresses the meanings of those concepts.

Aside: the statements which I focus on in this essay pick out things, and so they can be called *noun-centric statements*. Actions, processes and all other kinds of change over time, can also be picked out, of course, and something can be said about them, too. They can be called *verb-centric statements*.

Verb-centric statements have a more complex semantics than noun-centric ones,

⁵ For propositional attitudes, see (Chierchia & McConnell-Ginet, 1990; Chapters 5 & 6); (McKay and Nelson, 2014); (Schiffer, 1989); (Schwitzgebel, 2014). For Gricean implicatures, see (Chierchia & McConnell-Ginet, 1990; Chapter 4); (Schiffer, 1989). For a related discussion of both topics, see (Lyons, 1995; Chapters 8-12). The earliest account of Gricean implicatures, before that is what they were called, can be found in Nowell-Smith's *Ethics*. See especially his discussion of four rules of contextual implication, pp. 80-87.

as indicated by the thematic roles for things mentioned in them. However, if a semantic theory has to begin with one or the other, it is noun-centric statements it must begin with, because verbs are about what happens to nouns; processes are what happens to things.

Moreover, processes and other changes change the *states of things*, which are the sets of properties those things have at various points in time. And noun-centric statements are statements about states of things, with basic noun-centric statements ascribing one property to an object, the possession of that property being a state of that thing. Hereafter, unless otherwise indicated, "statement" will refer to a noun-centric basic statement.

For example, here is a statement:

(S1a) That rock is heavy.

In predicate logic, (S1a) is an instantiation of an existentially-quantified statement, which would be expressed this way:

(S1b) $\exists x(Rx \ \& \ Hx)$, (to be read as "There exists an x such that x is a rock and x is heavy".)

So a statement is a declarative sentence that does two things. It picks something out, and it says something about it. Its *subject term* does the first job; its *predicate term* does the second. In (S1a), "That rock" is the subject term, and it picks out an instance of a kind, the kind [Rock]. "is heavy" is the predicate term, and it ascribes a feature to the rock, saying that it has the property of being heavy, thus instantiating the feature [Heaviness]. So things are instantiations of kinds, and properties of things are instantiations of

features.

What the subject term picks out is the *referent* of the statement. What the predicate term picks out is an instantiation, in that referent, of a *feature*, that instantiation being a property of the referent (or a relationship it has to another referent).

The schema of a basic statement is shown in Figure 3.

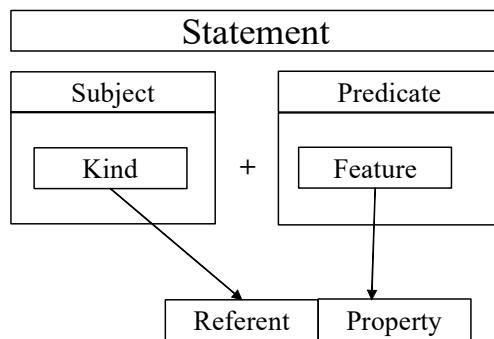


Figure 3. The Schema of a Basic Statement.

But the referents of many statements are not objects. Rather, they are concepts. For example, consider this statement:

(S2a) Rocks are heavy.

(S2a) is the background theory way of expressing the following predicate logic universally-quantified statement:

(S2b) $\forall x(Rx \supset Hx)$, (to be read as "For all x, if x is a rock, then x is heavy".)

(S2) has the kind [Rock] as its referent, and it ascribes the feature [Heaviness] to that

kind. In making or assenting to a universalized statement about a concept, a person expresses his belief that every instance of the subject concept is also an instance of the predicate concept. Since (S2) expresses the belief that every instance of the kind [Rock] is also an instance of the feature [Heaviness], any counterexample which remains a counterexample after attempts to discredit it, will force the person to retract his assertion of or assent to (S2), and admit that his belief was mistaken.

So we have two kinds of basic statement. *Object* statements are about particular things; they are instantiations of existentially-quantified statements. *Concept* statements are about categories of things; they are universally-quantified statements, statements about all particulars (if any) which fall under their subject term.

In set theory, there is a distinction between set membership and set inclusion. Set membership is a relationship between sets and things. Set inclusion is a relationship between sets and sets. Something is a member of a set just in case it satisfies the membership criterion for that set. A set is included in another set just case all members of the former set are also members of the latter set. Object statements express set membership; concept statements express set inclusion.⁶

In both object and concept statements, we have one thing being referred to, and two concepts, one concept indicating the kind of thing the referent is, and the other concept ascribing a property to that referent. For object statements, an explicit or implicit indexical is also part of the subject term. This is needed to pick out the particular object being referred to which is an instance of the statement's kind. For concept statements, indexicals are neither possible nor necessary. But a concept which is a referent in a statement is also an instance of a kind, that kind being a superordinate concept. However,

⁶ Here and later on, I will find it convenient to speak of concepts in terms of sets. But it should be kept in mind that these are the special kind of sets I have called categories. These are sets, unlike the mathematical ones, that are both time-variant, and associated with membership criteria that are often fuzzy.

because such relationships are generally widely-known, it is common for concept statements to leave the referent's superordinate concept unexpressed. We may say "Dogs are good pets", but we don't say "Animals which are dogs are good pets".

So statements are the nexus within which meaning, which is about universals, and reference, which is about particulars, interact to express information. This conjunction of two concepts and one referent is the structure of the association, within statements, of meaning and reference. Syntax, as I view it, is a means of expressing multiple basic statements in less time than it would take to list those basic statements, thus making it possible -- to anticipate my later neurophysiological analysis -- to keep more information in short-term memory.

Statement Tokens and Statement Awareness Events

A statement token is an instance of a statement, created by a specific person on a specific occasion. A statement token which did not conform to the semantics governing its concepts and their combination within that statement would, to the extent of its deviation from widely-accepted rules, be unintelligible.

But those rules themselves are derivative. They may be pedagogical; but they are not foundational. Dictionaries are periodically revised in order to "catch up" to changing patterns in how words are used. Etymological dictionaries show us changes in the meanings of words on, approximately, a time-scale of centuries. Revisions of standard dictionaries show us changes in the meanings of words on, approximately, a time-scale of decades. Conversations, in which speakers mutually influence one another's linguistic behavior, can show us changes in the meanings of words on a time-scale of minutes.

Note also that when the referent of a statement token is a particular object (rather than a

concept), the phonetic sequence or orthographic string does not uniquely identify the statement it is a token of. For example, the token "That rock is heavy" is a token of one statement when it refers to that rock over there, but is a token of a different statement when it refers to this rock over here. What makes the statements different is that they convey different information, being about different things.

When I am present on the occasion of your saying "That rock is heavy", and I hear and understand you, one statement token has been produced. But in addition there are two *statement awareness events*, yours and mine. Patterns of neural activity, in statement awareness events, are what give rise to patterns of linguistic behavior. Those patterns both conform to and determine the meaning and reference of statements. Now let's see what kind of behavior that is.

Section 2. Meaning and Reference

The meaning and reference of statement types and concept types are a function of their corresponding tokens, i.e. of the production or comprehension, on individual occasions by individual language users, of those statements and concepts. Over a community of language users, and a timeframe extending into the not-too-distant past, our referential use of concepts, and our co-ascriptional use of pairs of concepts, manage to converge to an extent sufficient to enable successful communication (as judged by participants in those communications) to take place. Let's see how this works.

Meaning: Concept-to-Concept Links

I will begin by discussing statements and concepts as types, and only after that turn to their physical realization as tokens.

In our background theory, the meaning of a concept is said to be expressed in its definition. For example, the meaning of the word "bachelor" is expressed in the definition "an unmarried adult male human being". How does this work out in the semantic theory being developed here?

To begin with, this definition summarizes a set of *entailments*, statements which are said to be true by virtue of the meanings of the co-occurring concepts, and which are taken to be consequently necessarily true. Such statements are called *analytic* statements. (Quine, 1961); (Rey, 2018).

(S3) is an entailment set for the concept [Bachelor].

(S3)

- (S3a) if x is a bachelor, then x is a human being;
- (S3b) if x is a bachelor, then x is an adult;
- (S3c) if x is a bachelor, then x is male;
- (S3d) if x is a bachelor, then x is unmarried.

As an alternative, we might first define "man":

- (S4)
 - (S4a) if x is a man, then x is a human being;
 - (S4b) if x is a man, then x is an adult;
 - (S4c) if x is a man, then x is male.

And then define "bachelor" like this:

- (S5)
 - (S5a) if x is a bachelor, then x is a man (equivalently, "All bachelors are men");
 - (S5b) if x is a bachelor, then x is unmarried (equivalently, "All bachelors are unmarried").

Each of (S3), (S4) and (S5) is a set of statements. Each statement picks something out (with its subject/antecedent term) and says something about it (with its predicate/consequent term). In addition, each of these statements are entailments. That is, they don't just happen to be true. They are *necessarily true*, true because of what their component concepts mean (and because of what the if...then and all...are statement schemas mean). Each entailment relates a pair of concepts, e.g. [Bachelor] and [Unmarried], [Bachelor] and [Male], [Bachelor] and [Adult], etc.

As I said earlier, and as this example illustrates, the referent of a statement need not be a particular thing. It may be a concept, in which case the statement ascribes a predicate concept to its subject concept. In making a concept-to-concept co-ascription, the statement implicitly asserts that every particular thing which is an instance of the subject concept possesses a property which is an instance of the predicate concept; equivalently, it expresses an inclusion relationship between the two concept sets.

Such statements are universally-quantified statements. This is illustrated in Figure 4.

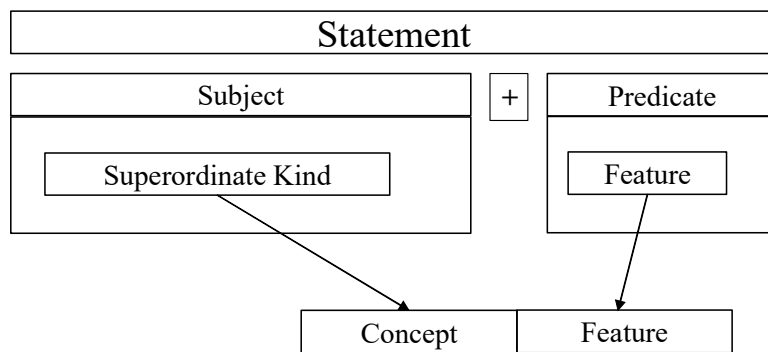


Figure 4. The Schema of a Concept Statement.

For example, consider statement (S5b), usually rendered as "All bachelors are unmarried".

[Bachelor] is a kind; [Unmarried] is a feature; "all" marks the statement as having no exceptions, and thus as being universally-quantified.

Most people, on most occasions, will agree that (S5b) is true. Indeed, this is a paradigmatic example of an *analytic statement*, one which is true by virtue of the meanings of its constituent concepts -- in this case, that [Unmarried] is part of the meaning of [Bachelor]. But sometimes, some people will not agree. Someone who

dissents from (S5b) believes that there is at least one counterexample to it, at least one bachelor who is married. So he produces what he thinks is a counterexample:

(S6) That man over there is a bachelor, but he's married.

However, we may challenge our dissenter, and the challenge will take the form of arguing that, in statement (S6), either the kind associated with that object over there, or the feature ascribed to it, is a mistake. To show that the counterexample is a mistake is to show (i) that the referent is not of the indicated kind [Man], and/or (ii) that the referent does not have the indicated feature [Married].

With many universally-quantified statements, however, those who believe they are true may eventually get tired of making specific arguments against specific counterexamples. After all, if someone is sure that all bachelors are unmarried, he doesn't need to discover that each one is, and won't feel inclined to argue that each one is. So his position can easily become "It's not just that there *aren't* any counterexamples to 'All bachelors are unmarried'. There *couldn't* be any. If anyone is a bachelor, he *must* be unmarried." The statement isn't contingently true, true because the world happens to contain no married bachelors. Rather, the statement is *necessarily* true.

But what makes a statement necessarily true? To say there *couldn't be* exceptions to a universally-quantified statement, and that any instance of its subject term *must* also instantiate its predicate term, expresses a very strong belief in the truth of the statement. But what accounts for these feelings of couldn't-ness and must-ness?

The notion of *possible worlds* in logic is used to clarify the modal distinction between necessary truth and contingent truth. (Hughes & Cresswell, 1968); (van Benthem, 2018). The idea is that a necessarily true statement is one which is true in all possible worlds,

while a possibly true statement is one which is true in at least one possible world (this real world, for example). A possible world is some situation described by a consistent set of statements, including all their deductive consequences. In a different but not unfamiliar use, one which gives a background theory feeling of comfort with this formidably technical piece of modal logic, a possible world is what someone, on some occasion, could imagine being the case. I will refer to these imagined possible worlds (which can also be situations in this actual world) as *scenarios*.

Of course, the possible worlds that a speaker, on an occasion, could imagine need not be restricted to scenarios that others would consider likely, and may extend to scenarios that others would consider so strange as to be almost unintelligible. Also, what someone could imagine on a given occasion depends on how strange a possible world he could conjure up, on that occasion. Today (perhaps after having read a science fiction novel) he might imagine possible worlds he didn't or couldn't imagine last week.

But there is no way to filter out plausible from implausible possible worlds (other than logical consistency). The basic empirical fact of being able to imagine counterexamples to a co-ascription of a subject concept and a predicate concept is not constrained by the judgments of others. As John Lyons said: "There is nothing in the structure of English which commits us to the denial of unfashionable or eccentric ontological assumptions." (1995; p. 142). Such commitments only come, Lyons continued, from "the notion of context-independent literal meaning with which many formal semanticists operate", and "is tacitly associated with their own context-dependent, philosophically challengeable, ontological assumptions." (ibid; p. 282.)

On the contrary, Lyons said:

"As always, reference is context-dependent: it is determined, first of all, by the speaker's general ontological beliefs and assumptions and, then, by his or her

more specific background beliefs and assumptions relevant to the particular context of utterance and often acquired in the course of the particular conversation to which the utterance contributes and of which it constitutes a part." (ibid; p. 188.)⁷

The description of any scenario, we may assume, consists of a set of statements. In the case of universalized concept statements, the speaker expresses a belief that there are no objects -- physical, mental or abstract -- which instantiate the subject concept but do not instantiate the predicate concept. The expression of her belief manifests a *co-ascriptional disposition*, specifically a disposition to not ascribe a subject concept to an object and, at the same time, the negation of the predicate concept to that same object.⁸ This disposition manifests itself, over multiple occasions of usage, as a pattern of neither producing nor countenancing the co-ascription, to the same referent, of the subject concept and the negation of the predicate concept. Such co-ascriptional dispositions create an extensive web of constraints on the concepts we are willing to co-ascribe in statements we believe to be true. Such dispositions constitute the roots of meaning.

Meaning Sets

We began with the background theory notion of the meaning of a concept, and developed that intuition at the level of analytic-philosophical and lexical semantics theories as the notion of a set of entailments which express the definition of a concept which, in turn,

⁷ Lyons is speaking about reference in these passages, but his point applies equally to meaning. (S5b), the statement challenged by the counterexample of (S6), *must*, it was claimed, be true, and thus is supposedly true in all possible worlds, and thus true no matter what this world is like or could be imagined to be like, and thus necessarily true. And as a statement whose truth is supposedly independent of both any actual and any possible scenario, the conclusion is that the necessity of the statement is based on its meaning, and on the meaning of its constituent concepts. This, by the way, is not a novel conclusion of the argument of this essay; indeed, it is an almost analytic-philosophical black letter conclusion.

⁸ This is the pattern of co-ascriptive behavior which, in 1972, I called "a disposition to disallow as misdescribed", a "DDM". (Johnston, 2015.) It is the disposition to avoid the co-ascriptional pattern I have just described which is, equivalently, a disposition to ascribe, or to always allow the ascription of, the predicate concept to whatever object to which one has already ascribed the subject concept. And this is, clearly, a disposition to treat an "All A are B" statement as analytically true, i.e. true by meaning.

expresses the meaning of that concept. Each statement in the entailment set for a concept is an analytic statement, one taken to be necessarily true.

But all we found, to explain the necessity of such statements, is that few fluent speakers would accept a description of a scenario in which there is claimed to be something which instantiates the subject concept but which does not also instantiate the predicate concept. Such speakers have, as I noted above, a "disposition to disallow as misdescribed" the ascription of a subject concept together with the negation of the predicate concept to the same object, on the same occasion. From Lyons' perspective, focusing on reference, such speakers have a disposition to believe that there neither are nor could be any such objects.

It follows that the necessity of these statements is not an all-or-nothing phenomenon. It is a matter of how thoroughly a co-ascriptional pattern has permeated the linguistic behavior of a community of speakers during some period of time, and the degree to which it has become part of the formal and informal sense of what those concepts mean.

It further follows that talk of entailment sets, in the strict, logician's sense, is misleading. A statement in an entailment set is a generalization over its statement tokens, produced by various speakers at various times; and for most concepts, exceptions to such generalizations will be tolerated, especially in conversational contexts. And, as the histories of languages prove, changes in dictionary definitions, and thus in entailment sets, do occur over time and sets of speakers. Clearly, these changes are gradual, so as they are going on (and they are always going on), the logician's notion of strict entailment has no place, except perhaps as something analogous to the concept of a limit in calculus.

Consequently, unless I am presenting descriptions of statement and concept *types*, and attempting to explain the associated beliefs in strictly analytic statements, and in definitions as sets of necessary and sufficient conditions describing fixed boundaries for

the meanings of words -- that is, unless I am speaking in the analytic-philosophical and linguistic science language in which such things are talked about -- I will no longer talk about entailments and entailment sets. Instead, I will talk about *inferences* and *meaning sets*. Each inference in the meaning set for a concept A is an inference to another concept B, and our belief that some of these inferences are necessary is an over-generalization of the fact that some of them are ones which most of us, on most occasions, would be strongly reluctant to admit exceptions to, those exceptions being things which are said to be instantiations of A but not of B. Since this reluctance is, for each of us on each occasion, a matter of degree, these sets expressive of the meaning of concepts are dynamic *graded* sets. Since type-level definitions of the meaning of concepts are abstractions from the co-ascriptive linguistic behavior of presumed fluent native speakers, on relatively recent occasions, further processed by lexicographic intuition about which inferences do and do not belong to that subset of inferences which should be organized into dictionary definitions, these formal dictionary definitions are doubly graded -- once at the individual token level, and then again as generalizations and over-generalizations across fluent speakers and across recent time. These doubly-graded definitions are then formalized and published, and presented to learner speakers as prescriptive for their use of those concepts.

Language is much more fluid and fuzzy-edged than it appears to be when we focus our attention on types which, eschewing a Mentalese or Platonic view of language, are reifications of patterns of linguistic behavior.⁹ Saussure, in disparaging diachronic semantics and focusing on synchronic phenomena, recognized only half of the interdependency between rules and regularities. Language does not exhibit just a behavior-on-rules dependency, like the dependency of chess games on the rules of chess.

⁹ One alternative, in the literature, is to take the meanings of concepts to be mental objects, and definitions of concepts as more or less accurate descriptions of these mental objects. Another alternative is to take the meanings of concepts to be abstract (mind-independent) objects, and definitions, again, as more or less accurate descriptions of them. In spite of extensive literatures, with impressive pedigrees, I assume, in this essay, that both alternatives are dead ends.

It also exhibits a rules-on-behavior dependency, the dependency of semantic rules on co-ascriptional semantic behavior. Absent this latter dependency, periodic dictionary revision would be unmotivated.

At the cognitive psychology level, a meaning set is a set of learned associations between the physical words representing the first and second concepts in a co-ascribed ordered pair of concepts. What is learned is a disposition to use, or to accept the use of, the second concept as ascribed to an object, on an occasion on which the first concept has already been ascribed to that object. At the cognitive sociology level, our linguistic behavior results in part from the constraints within our individual meaning sets, and then, in communicative acts, on the mutual influence the meaning sets of each of us in the language community have on the meaning sets of others, until we achieve a satisfactory coordination, one which enables most of us to think, most of the time, that we understand one another's statements.

Given that the analyticity of statements as types is an artifact of formal and informal lexicographic generalizations over patterns of co-ascriptive use, it follows, as Quine famously argued, that analyticity is a matter of degree. That is, it follows that meaning sets are graded. I have just argued that even for statement tokens -- occasions on which statements are produced -- and statement awareness events -- occasions on which statements are understood -- analyticity is a matter of degree, the degree to which the person producing or understanding the inference is willing to say that there might or might not be an exception to that co-ascriptive pattern, some possible scenario in which something is said to exist or is denied to exist, that instantiates the first concept but does not instantiate the second concept.

So entailment sets, and the strict definitions they express, are social artifacts generalized from patterns of co-ascriptive concept use plus lexicographic intuition used to draw a line

between inferences that are strong enough to be categorized as entailments, and those that are not. "Reverse engineering", as it were, from this dictionary-revision lexicographic output (and its informal analogues), we see that instead of the artifice of entailment sets, co-ascriptive behavior establishes *meaning sets* between ordered pairs of concepts -- the first concept being the definiendum and the second concept being part of its definiens, the first concept being considered as a kind, expressed in the subject term which picks something out, and the second concept being considered as a feature, expressed in the predicate term which says something about it.

The relationship between meaning and analyticity will be quite familiar to analytic philosophers of language. Equally familiar will be Quine's objections, most famously in "Two Dogmas of Empiricism", to using either one to explain the other. What we need, Quine argued, is an empirically-grounded account of one or the other. In *Word and Object*, he attempted to provide such an account for analyticity; and he admittedly failed, calling his concept of stimulus analyticity an at best *ersatz* description of the standard concept of analyticity, the one philosophers were (and still are) trying to explain. (1970; p. 66). Here, I am taking the other approach, attempting to provide an empirically-grounded account of meaning. But that account, I believe, cannot be completed without providing a similarly grounded account of reference, and of how the two interact in the linguistic behavior of fluent speakers. Those are the next topics.

Reference: Concept-to-Object Links

At the background theory level, the reference of a concept is whatever it is that it correctly refers to. Equivalently, it is the set of all and only those things that "fall under" that concept. For example, "horse" correctly refers to all and only those animals which are horses. Each such animal is a referent of that concept; each one "falls under" that concept (or "falls into" that category).

Concrete concepts are ones which refer to physical objects and their properties or relationships. Let's focus on these concepts first. So how does this happen? The notion that we notice different kinds of things, and also their properties, and then give names to them, is reminiscent of the story of Adam creating an ur-language, and as an account that science can develop, is deserving of no more consideration than is that story about Adam. For one thing, it is established neuroscience fact that a great deal of multi-level pattern recognition and consolidation occurs before the awareness of anything takes place, even something as simple as a sound or a color. The consolidation is even richer for the multi-modal perception of objects, such as horses or rocks.

But perhaps the ur-language story is not inconsistent with these facts. Perhaps all we need to do to rescue the ur-language account is accept that after sensory input is integrated sufficiently, Adam became aware of colors and sounds, horses and rocks, and that it was only then that he began giving names to them.¹⁰ Certainly, the not-yet-discredited notion that "natural kinds" really exist in the world around us, and maybe even social kinds as well, is at least consistent with the ur-language theory, and perhaps even supports it (in a dressed-up version, of course).¹¹

However dressed up, though, the ur-language story, and its closely-associated picture theory of meaning, is wrong. It has little of substance to say about the prototype feature of many concrete concepts, e.g. that apples will generally be more quickly and assuredly identified as fruit than, say, pomegranates. And it has nothing to say about the way that the boundaries of concrete concepts shade into one another, e.g. that we cannot always

¹⁰ This adaptation and simplification of the original Kantian account of the relationship between things and what we say about them is increasingly widespread in the literature, although its origins with Kant are often not noted. See, for example, (Jackendoff, 1988) and (Lakoff, 1988).

¹¹ For natural kinds, see (Brzović, 2018) and (Bird & Tobin, 2018). For natural and social kinds, see (Burge, 2010), (Smith, 2015), (Brown, 2016), (Johnston, 2016).

agree whether to say that a small container holding liquid is a cup or a mug, or that a stretch of water out on the lake is blue or green.

As for concepts with apparently precise physical referents, they need to be understood as limiting cases of concrete concepts in general, just as analytic statements need to be understood as limiting cases of strong inferences in meaning sets. Here's how that works.

At the cognitive psychology level of explanation, reference is a learned association of a physical word with sensory input or with the multi-modal integration of a set of sensory inputs. Each input, whether single- or multi-modal, results in a *gestalt*, more precisely a *perceptual gestalt*. At this level of explanation, the perceptual gestalt of a referent is an *image* of that referent. At the neurophysiological level of explanation, as we will see in Part 2, it is a group of neurons that will fire synchronously in response to the recognition of an image.

A further sensory integration links an image with another perceptual gestalt, the latter being an image of a physical (written or pronounced) lexeme. I will call this latter image a *wordform*, and will hereafter restrict the use of the word "image" to non-wordform images. The association of an image and a wordform constitutes a *lexicalized image*.

Lexicalized images are what cognitive scientists generally call *concrete concepts*. They are concepts that, at the background theory level of explanation, seem to name kinds of things and kinds of features that we can just see are out there in the world -- horses, rocks, green things, square things, and so on. They are concepts whose use is more strongly influenced by sensory experience -- whether currently present, remembered, imagined, or anticipated -- than by the inferences in the meaning sets of those concepts. *Abstract concepts*, by contrast, are the other ones. But the notion of a completely concrete concept or a completely abstract one -- a concept which is not influenced at all by either

meaning or reference (respectively) -- is not part of the semantic theory being developed here. That notion, I believe, is a generalization too far.

At the cognitive psychology level of explanation, a concrete concept is a disposition to assent to the ascription of the concept's wordform when in the presence of a corresponding image, a disposition in which either the wordform or the image can bring the other into conscious awareness.

The concrete concept of an apple, for example, associates an image of an apple with the wordform "apple". The *reference set* of the concept [Apple] is then, via the wordform "apple", the set of all those images that would elicit assent to the ascription "apple". In the case of this particular example, the reference set may remain relatively precise and relatively stable over an adult speaker's lifetime, and even over the community of fluent speakers of English during decades or more of use. In the case of many other examples, reference sets will be far more imprecise and far more interpersonally and temporally variable.

If challenged, we will often defend the *reference* of a subject or predicate concept by appealing to its *meaning* set. This is because the meaning set expresses the category membership criterion for that concept. However much it may seem so, one doesn't just "look and see" that the object is of the specified kind, because being of that kind means that it shares with all its companion objects in the reference set for that concept all the other features associated with it by its meaning set, as evidenced by the already-described co-ascriptive behavior of the kind concept and its inferentially-associated feature concepts.

Indeed, the subjective experience would commonly be reported as an experience of just seeing what kind of thing something is (or what kind of feature it has). But the fact of the

matter is that every concept has a meaning set, a graded set of inferential relationships to other concepts; and so a referent to which one or more of those inferentially-related concepts does not apply is, proportional to the strength of that inferential relationship, an invalid referent. Anything which is not a mammal is not a horse, no matter how closely it may resemble a horse.

As shown in Figure 5, the subject term of an object statement picks out its referent by doing two things. First, it indicates what kind of thing the referent is. Second, it uses an indexical to designate a specific instance of that kind, that instance being the referent of the statement. Consider once again (S1), the statement "That rock is heavy". In the subject term of (S1), "rock" does the first job; "that" does the second.

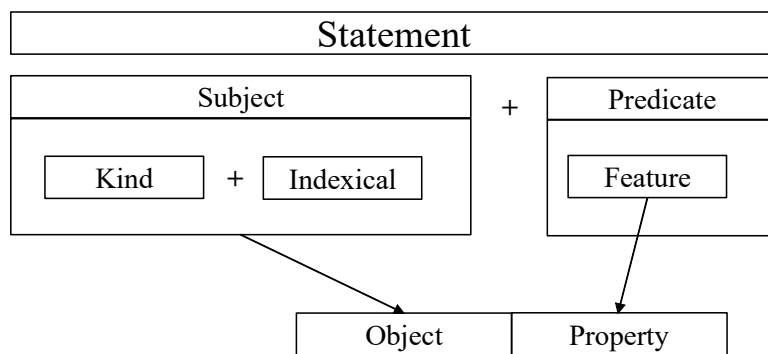


Figure 5. The Schema of an Object Statement.

In a statement about concrete objects, reference succeeds when and only when a speaker and a listener agree on what kind of thing the referent of the statement is, and on what property it possesses. In each case, the referential use of a concrete concept has been agreed on by speaker and listener. In the subject term of a statement, that concrete concept is a kind. In the predicate term of the statement, the concrete concept is a feature.

Concrete concepts are neurally realized as Hebbian-learned associations between a

hierarchy of neural recognition and integration events resulting in an image, and a hierarchy of neural recognition and integration events resulting in a wordform.

In addition, it will be useful to distinguish between a concrete concept as it exists in one person's consciousness at one point in time -- in the person's short-term memory (STM) -- and that concrete concept as stored in that person's long-term memory (LTM). The distinction is important because recognition involves the matching of the former to the latter. So when I need to emphasize the distinction, I will restrict the term *concrete concept* to refer to the current awareness / STM lexicalized image, and use *concrete engram* to refer to the amalgamated LTM memory of a history of concrete concept experiences with which the same wordform has been associated. All awareness of physical things and features, as we will see, is a matching of current images with engrams, a graded matching in which, on each occasion, a current image modifies an engram by being more or less accepted as an instance of that engram, or more or less excluded as an instance of it.

So for any but the most abstract concepts, just as we have graded meaning sets, we also have a graded set of lexicalized images. But the concepts which these lexicalized images belong to are not isolated. As Saussure emphasized, no concepts are. Concrete concepts are themselves nodes in a structuralist network of concept-to-concept links. These links constitute the meaning sets for those concepts. They are the inferences in those meaning sets.

We now have both reference and meaning as links in networks, the one linking concepts to things and properties of things, and the other linking concepts to concepts, each one part of a consolidation into one of these two networks.

Reference Sets

At the background theory level depicted in Figure 2, here's how we talk about the world. We see a horse and, when asked, say that it is a horse. That is, we use the word "horse" to refer to it. Or someone asks us for a cup of coffee. Knowing what the word "coffee" refers to, we go and get it for them. Given some thing, we can often say what kind of thing it is. Given a word, we can often pick out the things it refers to.

This is a deep background theory belief about language and the world. Concrete words correspond to things. Dressed up, it presents itself as a correspondence theory of truth, and a picture theory of meaning. In the heyday of logical positivism, the early Wittgenstein, he of the *Tractatus*, gave it its last gasp of intellectual respectability (which he shortly thereafter began to demolish, in the *Blue and Brown Books* and the *Philosophical Investigations*). The "Meaning ain't in the head" externalism of Putnam (1975), and the causal theory of meaning of Kripke (1980), are the unfortunate lingering remains of this not-yet-discredited view of how language relates to the world.¹²

For each of us, each occasion on which a concept (via its wordform) is paired with a physical object or feature (via its image) strengthens that association. Each occasion on which a concept is denied to correctly pick out an object or a feature, weakens that association. We learn to *see* the same kinds of things and properties and relationships of those things that others see by mutually influencing one another's (bidirectional) wordform-to-image referential associations, i.e. one another's concrete concepts. And at both the individual and societal levels, occasions of pairing or refusing to pair an image and a wordform extends or contracts the referential range of the concrete concept associating them.¹³ That concrete concept exists, over time, for each of us, as a

¹² On the long-standing, ongoing debate about externalism and internalism, in which the views of Kripke and Putnam are deeply involved, see (Lau and Deutsch, 2016); (Smith, 2015); (Johnston, 2016).

¹³ "refusing to pair" and "contracts" allude to the notion of an "anti-connection". I believe that a full account of the reference/meaning dynamism which determines the ever-varying and ever-evolving semantics of language includes not only graded connections, of both meaning and reference, but also

wordform-to-image engram, and each occasion of assenting to or dissenting from that wordform-to-image pairing modifies that engram.

Objectivity can only be established by an unobtainable God's eye point of view. It is a myth of background theory, albeit certainly an often useful one. But without that God's eye point of view, what we do have is an always dynamic, always under construction, *intersubjectivity*, mutually correlated and correlating learned patterns of associating a pair of sensory engrams -- an image with a wordform. Naturally, across time and persons, this correlation will be less than perfect; the reference sets for words, activated by sensory input, may vary over time for the same persons, and over persons at the same times. Together with meaning sets, the mutual influence of each of these person-relative sets on the sets of other persons, and their influences on ours, is the cognitive sociology explanation of the family resemblance phenomenon first noted by Wittgenstein, and of the prototype theories of meaning based on the work of Eleanor Rausch (1975).

At the cognitive psychology level of Figure 2, here's what's going on. In *reference* sets, concrete concepts are associated with images. The association is a bidirectional stimulus-response learned association of two sensory patterns, one for the wordform (via its own image) and one for the object or feature (via its image) it is associated with. Awareness of either can elicit awareness of the other.

In *meaning* sets, on the other hand, the association is a unidirectional stimulus-response learned association of two sensory patterns, each of the physical wordform of one of the concepts involved in the inference. The ascription of the first concept can prompt

graded anti-connections. Although it would require a lot of detailed work to include anti-connections in this material -- work I won't include in this particular essay -- it would essentially involve no theoretical advances, and could be expressed by copying this essay and changing inferences to their negations, "included in" to "excluded from", and so on. And this is not just a matter of dotting some *is* and crossing some *ts*. I believe that anti-connections are equipotent with their connectional counterparts in shaping the semantics of language.

acceptance of the ascription of the second concept, but not vice versa (unless the concepts are synonyms).¹⁴

At the cognitive sociology level, intersubjectivity is established by observing one another's pairings of wordforms with images, and of wordforms with other wordforms, and making mutually accommodating adjustments so as to reduce the occasions on which it seems that understanding one another is either imperfect or even absent.

So reference sets impose referential constraints on our linguistic behavior. But because of the mutual dependency of reference and meaning, I can't use "cow", for example, to refer to a horse because the definition of "cow" (its informally and formally accepted definition, i.e. its set of inferences strong enough to be accepted as entailments) includes concepts used to define "cow" which don't apply to horses.

But that's a matter of reference, you may say. And it is; cows just aren't horses. But in each of those cases, the refusal to accept, for example, that a real cow falls under the concept [Horse] is, in addition to referential constraints, because that concept has its own meaning set, and the concepts in that meaning set will also have constraints on them including, for example the inference from [Horse] to [Not having a prominent udder] and the correlative referential constraint of the concrete concept of an animal with a prominent udder will work against calling that particular cow, on that particular occasion,

¹⁴ To anticipate a neurophysiological discussion (in Part 2), the uni-directionality of inferential associations needs to be accounted for. Friedmann Pulvermuller's notion of neural sequence sets (2002) might be adapted for this purpose, although the order in an ordered concept pair {A,B} is not a temporal ordering, which is what sequence sets account for. In "Someone is unmarried if he is a bachelor", "unmarried temporally occurs before "bachelor" even though the inference is from "bachelor" to "unmarried". Nonetheless, Pulvermuller's work makes it clear that ordering is not an issue that is opaque to neuroscience. As a development of Pulvermuller's sequence sets, for example, we might start with the observation that between any two sets of communicating neurons, although each pathway is uni-directional, it usually turns out that pathways exist in both directions. So to explain the uni-directionality of inferences, we might speculate that inferentially co-ascriptive patterns are usually expressed with the first concept in the inference being presented first, and that this physical pattern would lead to the atrophy of neural connections in the opposite direction.

a horse.

Reference, seen in stimulus-reponse connections between wordforms and images, is constrained by the meaning sets of those wordforms. And the concepts in each of those meaning sets have their own reciprocally-influencing lexicalized images, and their own meaning set connections to other concepts.

There is a parallel here with mathematical sets. The things which are the members of a set must conform to the membership criterion for the set, and that criterion must designate all and only those things. The reference set of a concept is the (time- and person-varying) collection of all the things that fall under that concept. The concept corresponding to the set is the kind that those things are. The meaning set of a concept is the (time- and person-varying) membership criterion for the set. Wordform-to-image associations display the concept's reference set. Wordform-to-wordform associations display the concept's meaning set. Objects, via their images, are members of that set. Related concepts, via their inferences, are sets subsumed under that concept's set.

Indeterminate Reference: Association Sets

Existentially-quantified statements are "Some A are B" statements. Object statements are instantiations of these existential quantifications, picking out a particular and saying of it that it is both an A and a B. In the canonical basic statement form we are using, such a statement would be "That A is B", where A plays the role of a kind that the indicated referent belongs to, and B plays the role of a feature that it instantiates.

Besides the form "Some A are B", uninstantiated existential quantifications can also be expressed in other ways, such as "Many A are B" and "Few A are B". Such statements are like promissory notes; if challenged, it must be possible to find a referent which is of kind

A and which possesses property B. "Most cats have tails" can easily be confirmed by pointing to a sufficiently large number of the vast majority of cats which do have tails. "A few cats do not have tails" can be confirmed by pointing to a couple of Manx cats.

The concept pair [Cat] and [Having a Tail] are associated, but that association is not part of the meaning set for [Cat]; it is not an inference because it is not a universal quantification. Nonetheless, it is an association, a regularity in the world around us that we can rely on. So it seems that we can supplement the meaning sets of concepts with what I will call *association sets*. Each concept pair in an association set I will call a *generalization* because we can think of it as being formed by moving an inference from a concept to a superordinate concept, thus generalizing the inference and often, in the process, losing its universality. As an inference, the ordered pair $\{A_1, B\}$ represents the statement:

(S7a) $\forall x(A_1x \supset Bx)$, (to be read as "For all x, if x is A_1 , then x is B".)

But say that the concept A_0 is a superordinate concept, such that all A_1 -- and other kinds of A, but not all other kinds -- are A_0 . Then if we generalize the A_1 inference link and "copy it" to A_0 , we will have to drop the universal quantification. The result is a generalization in the association set for A_0 :

(S7b) $\exists x(A_0 \& Bx)$, (to be read as "There exists an x such that x is A_0 and X is B".)

Inferences at the weak end of meaning sets shade into generalizations at the strong end of association sets. If I'm not sure whether or not all A_1 are B, I might still be quite sure that some of them are, and that some other kinds of A are too. On the other hand, I may know quite well that all of the A_0 which are A_1 are B. In the first case, generalizations are a way

of hedging our bets; in the second case, they result from regarding the often tedious job of using a concept which represents only the members of a specific reference set as unduly onerous and also unnecessary; using a more general concept will often be enough to "get the point across", and may enhance communication with others who may lack fluency with the more precise concept.

Association sets obviously influence linguistic behavior. But unlike inferences, they are relationships among *unordered* pairs of concepts. If some A are B, then some B are A. Nonetheless, generalizations play their concept-to-concept part in the counterexample exchanges which are the paradigmatic form in which we influence the semantic behavior of others, mutually aligning our meaning and reference sets until, it is hoped, agreement can be reached on the validity of the counterexample, or of other counterexamples retreated to in the course of the conversation, or on the conclusion that there are no counterexamples.

This is the contribution which association sets make to semantics. They permit us to make conceptual distinctions which allow us to reject a counterexample by arguing that the correct concept for the referent of the counterexample is one whose wordform represents a superordinate concept in a conceptual hierarchy, that the counterexample does not show an A_1 which is not B, but only an instance of some other subtype of A_0 , in which case the counterexample is valid against the inference "All A_0 are B", but not against the inference "All A_1 are B".

Meaning and Reference: Recursion

Association sets also permit us to generalize already accepted inferences, moving up conceptual hierarchies to propose, for example, that not only all A_1 are B, but also that all A_0 are B. Association sets leave us free to search, among hierarchical sets of concepts, for

that highest-level concept at which an inference still holds, and the one above it at which the inference first fails. This is, in a way, what thinking is -- the search for descriptions which are so precise that they apply to all and only those things believed to be similar in some way, thus being formulable as universally-quantified statements. Existential quantifications, then, are what we must settle for when we can't be that precise.

So meaning exists not just to make reference possible, but also to allow us to generalize, to recursively apply category membership descriptions to concepts increasingly remote from sense experience and, conversely, to allow us to reach down from those remote concepts to the things they are indirectly about, those things being either the objects corresponding to strongly concrete concepts, or objects in the sense of concept reifications.¹⁵

In both cases, we are involved with a web of words, a web of the physical representations of lexemes, including those lexemes close enough to sensory experience to be concrete concepts. On the other hand, the physical existence of concepts remote from sensory experience is their existence as the sense objects which are their spoken or written wordforms. As such, no sensory experience other than the perception of their wordforms is relevant to them, except indirectly, through the aforementioned transitivity of the recursively-formed meaning sets of increasingly less abstract words.

¹⁵ To a philosophical realist, "concept reification" is just name-calling because corresponding to any concept there will be things that "fall under" that concept, and those things ipso facto exist. For realists, as Quine said, "to be is to be the value of a variable" or, in more traditional terms, to be is to be a particular which is an instance of a universal. To a philosophical nominalist, on the other hand, the only things that exist are the referents of strongly concrete concepts. An even stronger version of nominalism would exclude even strongly concrete concepts unless their images were said to represent objects -- things that exist "on their own", rather than properties of objects. I mention this because I think that a more fully developed commentary on ancient and still current discussions of realism and nominalism, from the perspective and in the terms of the semantic theory being developed here, might prove interesting.

Meaning and Reference: Unification

A universally-quantified statement expresses the claim that all instances of its subject term concept are also instances of its predicate term concept. The statement can be proven false if an instantiation of its subject concept can be found which is not also an instantiation of its predicate concept. But determining this involves reference. In the case of noun-referent statements, it involves reference to an object, i.e. to an instance of a kind in the case of the subject concept, and also reference to a property, i.e. to an instance of a feature in the case of the predicate concept. A counterexample is an object of that kind which lacks that property, something which is an instance of the indicated kind-concept, but not an instance of the indicated feature-concept.

We have seen, in the account of meaning and reference so-far given, that both are constraints within which linguistic behavior manifests itself. The constraint-conforming patterns of behavior are manifested, for all to see, in the statements we make, and manifested for us in the statements that others make. Over time, and across a webwork of linguistic communities, we converge on graded patterns of wordform-to-image pairings that constitute the referential use of those concepts, and we converge on graded patterns of co-ascription relationships among concepts that constitute their meaning. Meaning and reference, both anchored in the association of sensory patterns, constrain the use of concepts in statements, and determine which statements we will say are true and which false, which are semantically well-formed and which are not. Because the statements we make are made in conformity with our evolving assimilation of and influence on these evolving societal convergences, we can understand one another and exchange information by making statements.

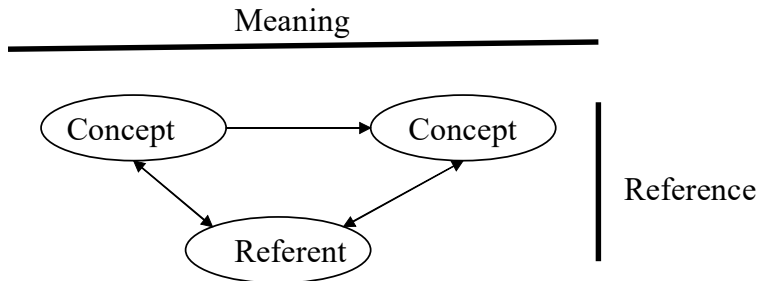


Figure 6. The Interplay of Meaning and Reference.

So in all statements, as shown in Figure 6, either an object or a concept of a particular kind is picked out, and by ascribing a feature to the concept, or a property to the object, something is said about it.

For object statements, if what we pick out with a kind term would not be counted as an instance of that kind by most others, or what we pick out with a feature term would not be counted as an instance of that feature by most others, then communicative reference has failed. We don't share the same concepts, even if we use the same wordforms.

Otherwise, we and those others must accept the statement as true, because we have, separately as it were, accepted that the object picked out is of the indicated kind (a ball, say) and the property picked out is an instance of the indicated feature (blue, say). Having accepted each of the references, made of the same object on the same occasion, we have also accepted both of them. The conjunction is not an added ingredient. We cannot fail but agree that that particular ball over there is blue.

Is this an instance in which reference alone determines the truth of a statement? No.

Concept-to-concept co-ascriptional constraints are also involved. We will not accept that the designated object is a ball if it doesn't meet the inferential criteria, and many of the associational criteria, for being a ball. For example, if a spherical object is very small, a shotgun pellet for example, we would generally hesitate to call it a ball. That is, in our

verbal behavior, we will usually not use "ball" to refer to a very small spherical object, and we won't because of that generalization about the size of things called balls.

So in statements, reference and meaning do not fail or succeed independently. Successful reference requires conformance to category membership criteria, i.e. to the meaning sets of the concepts, each of whose strongly inferential relationships to other concepts is itself a statement which must be true of the referent or feature referred to. A category membership criterion, when stated, is stated as a rule; but what it describes is a pattern of behavior. And the concepts in those category membership statements are related to a network of other concepts, each with their own co-ascriptive relationships, without which each of those second concepts would be semantically isolated, and thus meaningless.

Add an inference to a concept's meaning set, or remove an inference from it, and you've changed the meaning of the concept. Do that and you've changed its category membership criterion. Do that and you've changed the membership of the set of which that concept is the kind. Do that and you've changed the referential range of the concept.

Conversely, change the membership of a set and you may require changes in its category membership criterion. Do that and you've changed the meaning of the concept. Do that and you've aligned the meaning set of the concept to conform to its modified referential range.

Meaning and reference are inextricable and mutually-determining. I will call this the *meaning-reference-reciprocity thesis (MRRT)*. It is the most basic thesis of TM, this semantic theory about meaning and reference.

Section 3. Measuring Meaning, Reference, and Semantics

To facilitate experiments which can quantify the strength of reference and meaning as they influence ascriptive linguistic behavior, we need to focus separately on lexicalized images and on ordered lexical pairs, observing the influence of wordform-to-image associations and also of wordform-to-wordform associations. Then we can attempt to measure each component of semantics, and derivatively establish the graded distinction between true as a matter of fact and true as a matter of meaning.

To do all this, we need to set up test cases. Test cases for reference are, roughly, the well-known word-to-picture matching experiments used in cognitive science studies of language. Test cases for meaning are, roughly, the word association experiments also used in those studies.

Lexical Pairs and Canonical Forms

Current research into lexical semantics which focuses on lexical pair co-occurrences has correctly focused on these pairs as an object of study. That is indeed where meaning resides. But this net is cast too wide, since co-occurrences can reflect matters of fact as much as they do matters of meaning. I maintain that the atomic unit of the meaning component of semantics is to be found in the *co-ascriptio*n of ordered lexical pairs, i.e. their co-ascriptio

n to the same referent, the same thing. In their simplest form, the co-ascriptio

n of ordered lexical pairs occurs in statements, in the basic statements already described in which one referent is picked out, and one property or relationship ascribed to it. This requires us to translate the sentences selected from some corpus into the canonical form of basic statements, and I have already indicated how to do this. These canonical statements are the raw material from which

meaning can be extracted but, by themselves, they still do not distinguish between matters of fact and matters of meaning, between statements grounded in concrete concept ascriptions and those grounded in concept-to-concept co-ascriptions, or between statements expressing strong word-to-object associations and those expressing strong word-to-word associations.

If the concepts in the ordered pairs {Bachelor, Man}, {Bachelor, Adult} and {Bachelor, Unmarried} are related by meaning, as they are in their standard senses, then there just aren't any referents which are bachelors but not men, bachelors which are not adults, or bachelors which are not unmarried. Moreover, there couldn't be. To shift from an ontological to a language-behavioral description, any co-ascription of the first concept, [Bachelor], and the negation of the second concept, [~Man] or [~Adult] or [~Unmarried], would be considered a semantic mistake, one revealing that the speaker did not understand what one or both of that pair of concepts mean.

Notice that the co-occurrences of interest are not just lexical pairs; they are *ordered* lexical pairs. [Adult] can clearly be paired with many concepts which are mutually exclusive with [Bachelor], such as with [Woman]. The question is whether or not [Bachelor] can be paired with such concepts as [Woman] or [Child].

Since [Woman / Man], [Child / Adult] and [Married / Unmarried] are mutually exclusive, we should find few {Bachelor, Woman}, {Bachelor, Child} or {Bachelor, Married} co-ascriptions in our corpus. If even a single semantically acceptable one of these co-ascriptions can be found in data provided by fluent speakers, e.g. if some speakers express the belief that there is at least one woman bachelor, then the inference from [Bachelor] to [Man] is not necessarily true, true by virtue of the meaning of those two concepts. But if the inference is true by meaning, then there can be no exceptions to it. (Of course, this is strictly true for a speaker on an occasion, i.e. for that particular token

of that statement, and that statement awareness event. Speaking about the statement itself, i.e. the statement as a type, is to speak about an informal and formal lexicographic generalization over observed patterns of co-ascription.)

Nonetheless, each co-ascription of a pair of concepts is part of the meaning set or association set for the first concept, the one in the subject term of the statement.

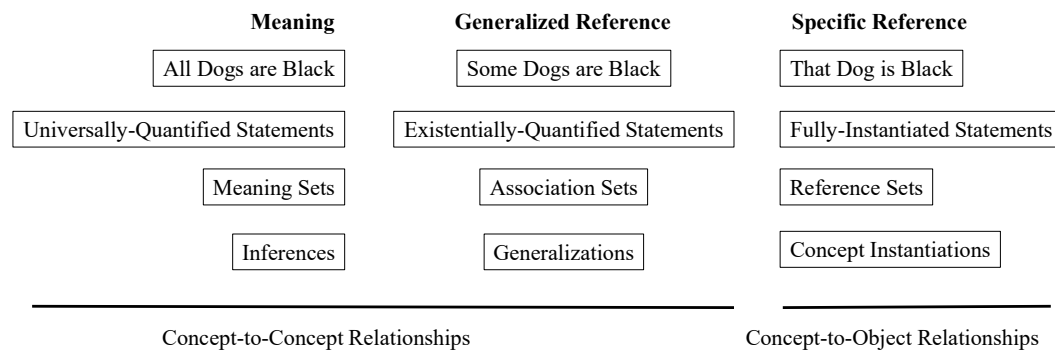


Figure 7. The Semantic Continuum.

As illustrated in Figure 7, meaning sets and reference sets lie along a continuum -- the *semantic continuum* -- with association sets lying between them. At one end of this meaning-reference continuum are concept-to-concept co-occurrences constrained by *meaning*, and to the extent that they are pervasive in a language community, they are expressed in analytic statements, ones necessarily true because nothing would be countenanced as an exception to them. Inferences in meaning sets express constraints on ordered pairs of concepts, and the strong inferences are the analytic statements making up the definition of the first concept in those ordered pairs. As a result, the logical form of inferences is that of universally-quantified statements, statements about all referents of the kind indicated by the subject term.

At the other end of this continuum are word-to-object co-occurrences constrained by

reference, existing as lexicalized images which establish the reference of the subject and predicate terms. These statements are, in the language of background theory, true (or false) as matters of fact. These statements claim only that a particular referent, one for which there is a corresponding lexicalized image, has the property ascribed to it by the second concept. Thus, the logical form of a statement at this end of the semantic continuum is that of an instantiation of an existential quantification.

Between meaning and reference are *associations* which claim, without specification, that there are some (or many, or few) instantiations of a pair of concepts. In the example in Figure 7, the generalization involves a pair of lexicalized images (i.e. concrete concepts). However, existential generalizations also exist between pairs of abstract concepts; for example, the statement that some prime numbers are less than one-hundred. Others combine lexicalized images with concepts which lack any direct association with images; for example, that some dogs are cute.

There is nothing more to the true as a matter of fact vs. true because of meaning distinction than this reference/association/meaning semantic continuum. Along this continuum, there are no dichotomies. Formal and informal definitions of terms are based on the relative infrequency of exceptions to ordered concept pairs, and on lexicographic intuitions about which pairs in a meaning set are so well-established that ascribing the first concept and the negation of the second concept to anything should be treated as a semantic mistake, as a description of no actual or possible thing at all. But in many cases, when we speak of the ascription of concepts to referents, we are speaking of ascriptions mediated by lexicalized images. And for some lexicalized images, the association of the wordform and the image seems a simple matter of seeing what is in front of you, while for others, we may feel doubt about whether or not the referent is correctly picked out by the wordform. Matters of fact are matters of degree just as much as are matters of meaning. Moreover, because only in limit cases (if there are any) will we exclude either

reference or inference as a determining factor in deciding the truth of a basic statement, there are, except for these limit cases, no statements true purely as a matter of fact, and no statements true purely as a matter of meaning.

When there are significant patterns in the linguistic behavior of co-ascribing a pair of concepts, that behavior will have neurophysiological origins, origins of what I have called linguistic dispositions, specifically, dispositions to disallow as misdescribed the co-ascription of the first concept and the negation of the second concept, and dispositions to disallow the denial of the association of a wordform and an image. So what we are really after -- and what statistical patterns in ascriptions and co-ascriptions reveal, when significant enough -- are these linguistic dispositions and their neurological foundations.

The "significant enough" qualification is necessary because our overt co-ascriptive behavior can be due to a wide range of factors, including speaking ironically, making a slip of the tongue, or speaking while in a drug-altered state. Is there any way to filter out this noise from genuine co-ascriptive behavior?

There is, but it requires us to move into the lab, and set up experiments with fluent, neurologically unimpaired speakers as subjects. Because in our pursuit of semantics, what we really want to know, for any statement token, is whether or not a speaker *would* countenance the co-ascription of a first concept and the negation of the second concept, or would countenance ascribing a kind concept or a feature concept to a particular referent. And this involves asking the speaker. Her linguistic behavior is on display in linguistic corpora; but her linguistic dispositions are not.

Measuring Meaning: Concrete Concepts

From a corpus, extract a set of universally-quantified basic statements. Let "All A are B"

be one of them; more formally, "for all x, if x is an A, then x is a B"; more concisely, the ordered concept pair $\{A,B\}$, in which it is asserted that every thing which is an instance of the kind A also instantiates the feature B.

We wish to analyze the response of a set of subjects to this inference. Each response, of course, is the response of a particular person, on a particular occasion, to a token utterance of the basic statement asserting that ordered pair.

So we have $\langle P, O, S_{\{A,B\}} \rangle$, $\langle \text{Person, Occasion, Statement} \rangle$. For this event triple, we ask the following question:

Q0: is there at least one instance of $\{A,B\}$ in the real world?

If the answer is No, then substitute $\sim B$ for B, and continue. If the answer is still No, discard the concept pair. It will have interest, but its interesting features are peripheral to our current concerns. Otherwise, we have a Yes answer. To avoid the awkward "B or $\sim B$ " which would otherwise intrude, I will just say that we have an affirmative response to the question of whether or not there is an instantiation of the ordered pair $\{A,B\}$ in the real world. After all, $\sim B$ is just as much a concept as B is.

Having a concept pair to test, we proceed to questions 1 and 2.

Q1: are there any exceptions to $\{A,B\}$ in the real world?

Q2: is an exception conceivable? (Or, does it make sense to talk about a non-B A?)

A (N/N) response to questions 1 and 2 marks $\{A,B\}$ -- for that person on that occasion -- as a strong inference. Grade it as CoA (co-ascriptional strength) =1. The $\{A,B\}$ pairs that

occur in definitions are strong inferences.

A (N/Y) response marks {A,B} as a universal generalization about the real world. Grade it as CoA=2.

A CoA=1 statement token is one true by meaning; a CoA=2 statement token is one true as a matter of fact. This is as close as empirical data will bring us to a clear dictionary/encyclopedia distinction.

For a (Y/0) response (with 0 indicating that question 2 is irrelevant), the subject indicates her belief that some As are B, and others are not. We have now dropped {A,B} as a universal quantification, and changed it to an existential one. We have moved from an inference to a generalization, and now ask:

Q3: would a non-B A be quite commonplace (Y), or quite exceptional (N)?

A (Y/0/N) response marks {A,B} as a strong generalization. Give it a CoA=3. A (Y/0/Y) response indicates a weak generalization. Give it a CoA=4.

But clearly, our co-ascriptive dispositions do not come in four discrete categories. They vary from strong inferences to weak generalizations, and do so along a continuum of disposition-to-disallow strength. For the same <P, S> pair, different occasions of asking these questions will often produce different answers. For the same <O, S> pair (with occasions identified with large-enough-sized time periods), different persons, on the same occasions, will often produce different answers.

So, for each {A,B} pair of concepts, we will have compiled a set of answer sequences, a set whose members are:

CoA 1:	N/N/0	All	Possible	Analytic	inference
CoA 2:	N/Y/0	All	Actual	Synthetic	inference
CoA 3:	Y/0/N	Some	Many	Strong	generalization
CoA 4:	Y/0/Y	Some	Few	Weak	generalization

Figure 8. Grading Co-Ascription Dispositions

So let a *meaning evaluation unit* for a pair of concepts have this form:

$$\langle P, O, S_{\{A,B\}}, CoA \rangle$$

This structure represents a Person, on an Occasion, making a co-ascriptional strength judgment about a Statement.

The next step, of course, is to sum each CoA ranking for each $\{A,B\}$. An average rank can then be calculated for (i) all its $\langle P, O \rangle$ data units; all its $\langle P_1, O \rangle$ data units for a given person P_1 ; all its $\langle P, O_1 \rangle$ data units for all persons on a given occasion O_1 . A classical diachronic analysis of language change over a given period of time would track the change in average rank, for each $\{A,B\}$ for $\langle P, O_{x,y} \rangle$, i.e. for all responses from all persons which took place during time period x,y . Over a sequence of time periods, we could then trace changes in the meaning set and the association set for concept A.¹⁶ We could see, not just that the concept did change its meaning, but also what components of its meaning set were dropped, and/or added, and/or changed co-ascriptional strength. And we can do the same for the association set for concept A.

¹⁶ What meaning sets and association sets have in common, because of which they are considered together, and separate from reference sets, is that meaning and association sets are both quantified, i.e. they contain variables. Reference sets do not. This is another way of pointing out that one component of semantics (meaning and association sets) is between pairs of concepts, while another component of semantics (reference sets) is between concepts and images.

We can also compile statistics for the meaning and association sets for B, and for all the other concepts in the meaning and association sets for A. And we can continue this process indefinitely. Ideally, if B is part of the meaning of A, and C is part of the meaning of B (and so on), then C (and so on) is also part of the meaning of A. Co-ascriptional strength is, ideally, a transitive relationship. However, in actual empirical fact, real human subjects are not perfectly rational, so when queried about the CoA strength for the pair {A,Q}, for which she has already indicated a CoA=1 for each of {A,B}, {B,C} ... {P,Q}, she may assign {A,Q} a CoA other than CoA=1.

Various analytical questions, whose answers I do not know, would include these ones.

- If the A to Q string of concept pairs are all CoA=1 except for one of them, which is CoA=3, would this make the {A,Q}, for a perfectly rational person, CoA=3?
- Or would it make it the average of the CoAs ($15*1 + 1*3 = 18/16 = 1.125$)?
- Or would it depend on where in the chain the CoA=3 appeared?

Whatever the answer is, if we could compare transitively inferred CoAs with directly tested CoAs, I suspect that the discrepancies, traced across a web of concept pairs, might prove worthy of further investigation.

Consider also, for example, our famous Kripkean example of a basic statement, "Water is H₂O". In some judgment events, it will receive a CoA=1, but in others a CoA=2. Can we explain this discrepancy and, in so doing, reveal the nature of analytic a posteriori statements?

Either ranking expresses the judgment that the statement is true. For those with a robust

knowledge of chemistry, and its relationship to physics, the set of other inferences in the meaning sets for "water" and for "H₂O" will be so mutually constraining that it will seem a misunderstanding of the concepts to suggest that there might be a counterexample. For those people, the judgment is likely to be CoA=1. For others, with little knowledge of chemistry, "Water is H₂O" will be accepted as a well-established fact; for them, the judgment is likely to be CoA=2.

It follows that the statement itself, the statement as a type, is neither true exclusively because of how the world is nor true exclusively because of the meaning of the words used in the statement. The matter-of-fact vs. matter-of-meaning distinction does not apply to statements; it applies to our judgments about those statements, made by each of us on specific occasions, and about the reasons we have for believing that true statements are true, and false ones false. And by forcing our dispositions to disallow counterexamples into four discrete categories, I have significantly oversimplified the actual co-ascriptional constraints that apply to a concept pair.

Note also that it is part of this semantic theory, of course, that concepts and statements, even as types, are not timeless, even though we talk about them, in our synchronic mode of speaking, in abstraction from speakers and times. Change in etymological dictionary time, and in standard dictionary revision time, make it clear that statements are not semantically stable, because the concepts used in them change their meanings over time. And I disdain any appeal to abstract objects such as propositions, in their role as the stable meaning shared by synonymous statements.

But as the MRRT tells us, for strongly concrete concepts, ones with robust reference sets, these judgments about ordered-concept-pair co-ascriptions are not independent of the reference sets of these concepts since co-ascription judgments are about the existence or non-existence of instantiations of those pairs, and of the possibility or impossibility of

using those pairs to describe anything at all, real or imagined. For object statements involving concrete concepts, their truth or falsity is constrained by lexicalized images linking their concepts to objects and so we are inclined to say of them that they they are true (when true) because that's just how the world is, and false (when false) because they fail to describe what they are purportedly about.

But that is almost never the complete story. The most concrete concepts have inferential relationships with other concepts. In situations in which it seems that we can just see that a particular object instantiates a concept, it is sometimes the case that the identification of that object as an instance of concept A is overruled by the claim that it "can't be an A", and that claim is expanded into the explanation that if that object were an A, it would also be a B, but it is not a B. And the force of the objection does not depend on whether or not B is itself a concrete concept.

Because of the MRRT, there is no dichotomy between matters of fact and matters of meaning, between statements true because they correctly describe what things *are* really like, and statements true because they correctly describe what things *must be* really like. And conversely, between statements false because they fail to describe what things are really like, and statements false because they fail to be descriptions of anything at all.

Measuring Meaning: Abstract Concepts

Before I turn to the measurement of reference, I will consider the situation for *abstract concepts*. The more abstract our concepts are, the farther removed they are from sensory experience, and so the more their combinations in statements are subject only to the constraint of inferential links among those concepts. A system of purely analytic inference links is a system to which the tools of deductive logic can be applied, tools with which we can conclude that concept pairs generated by these links are true if the

statements with which the chain of links begins are true.

For concept statements relating abstract concepts, a counterexample will not be a particular object that we can "just see" is an A which is not a B. Rather, counterexamples will be based on conflicts among meaning sets. This way of challenging a concept statement is to find an inference in the feature concept's meaning set that contradicts an inference in the kind concept's meaning set, i.e. that an inference "transitively down the line" from {A,B} brings up a contradiction. For example, to challenge the concept statement "All A are B", we might present as true the statements "All B are C" and "No C is A". If those two latter statements are strong inferences (with CoAs close to 1), then no matter what our senses tell us, we may be forced to conclude that not all As are B.

It is now time to turn to the measurement of reference, the constraint on our ascription of concepts which, together with meaning, constitutes the full semantics of those concepts.

Measuring Reference: Concrete Concepts

I begin with lexicalized images, i.e. concrete concepts. With them we can isolate *reference* just as, by focusing on ordered concept pairs, we isolated *meaning*.

In Figure 9, I show a relationship between a word and an object it picks out. Next, I show the statement awareness event in the mind of a speaker, in which the word and object have generated in consciousness a concrete concept consisting of a wordform and a perceptual image, each the result of pre-conscious processing taking place among sensory stimuli impinging on our sensory receptors. This awareness event exists in short-term memory (STM) and, without "rehearsal loops", would exist there for a few seconds. With cortico-thalamic rehearsal loops (which are fairly-well understood), it may exist in STM for a few minutes. But for that concrete concept to be accepted as a valid ascription, both

the image and its wordform must match the long-term memory (LTM) of each one closely enough to be accepted. Calling an apple an airplane would not result in a stable STM concrete concept event.

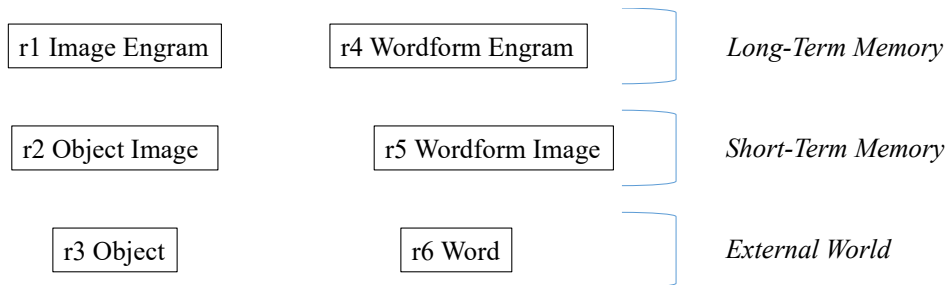


Figure 9. Measuring Reference.

The terms in each of the boxes in Figure 9 have been introduced previously. Their numerical sequence has no significance. The prefixes r1...r6 are merely labels for the boxes in Figure 9.

When an object is identified by a subject as a kind, perhaps from a list of kind terms, the task is an object-to-word match.¹⁷ The experiment should be set up so that there is no apparent distortion in the r3-r2 (or r6-r5) transition.

So our question is what causes an image to prompt assent to the ascription of particular word to that image.

The process begins with a present object (r3) producing an object image in a subject's

¹⁷ If a kind term is picked from a list, the match experiment is a "best fit" experiment. That kind of experiment measures the competition among lexicalized engrams as to which image best matches the image in current awareness. On the other hand, if a kind term is elicited without lists or suggestions, the match is a measurement of the correlation between the memory of a wordform/image pair and the current association of a wordform and an image. These are two related but significantly different empirical phenomena, which I believe would produce related but significantly different neurophysiological data.

brain (r2). This involves quite complex neural events, both within each sensory modality and across their integration, but those neural processes occur prior to conscious awareness and so I exclude them from consideration. Then awareness is focused on (r2), the object image.¹⁸

The process of matching current awareness to memory is the sequence of events r3-[r2-r1-r4-r5]-r6.¹⁹ At the cognitive psychology level, it is the sequence r2-r1-r4-r5. At that level, the image of an object (r2), in STM, is compared to one or more engram images (r1) from LTM. A match exists if and only if one of the potential engram images which match to (r2) is associated with a wordform engram (r4) which matches the STM wordform (r5). If it does, then the concrete concept engram, consisting of an image and a wordform (r1+r4), matches the current concept's image and wordform (r2+r5), which is the subject of awareness in STM. The object of which one is currently aware is judged to be of the kind ascribed to it (r3+r6).²⁰

The match of an object to a word, and of a concrete concept in STM to one in LTM, is a

¹⁸ As I use the term here, awareness (specifically, perceptual awareness) is the focus of attention on an object image. Although background theory speaks of awareness of an object, not of its brain-internal image, I use the term in this specific neuroscience sense in this discussion.

¹⁹ So in the string r3-[r2-r1-r4-r5]-r6, the pair outside the brackets relate the cognitive psychology level description inside the brackets to background theory.

²⁰ A philosophical issue is the correct semantic relationship between the levels of description shown in Figure 2. I have been relating these levels of description in this essay. But it must also be pointed out that objects and processes from any one level are not legitimate as part of the explanations in any other level. For example, mental objects (e.g. ideas) and processes (e.g. thinking, believing, etc.), from the background theory of mind level, cannot function as elements in a neurophysiological explanation of neurons and neural processes -- except as placeholders, to be filled in later.

Each level, in terms of the explanations within it, is a closed hermeneutic circle. And yet there are, and must be, semantic correlations between these hermeneutic circles. So *explanations* are valid only within hermeneutic circles. Across hermeneutic circles, we can provide only *correlations*. But if we don't understand how to distinguish explanations from correlations (and cognitive psychology and neuroscience literature makes it clear that it is sometimes difficult to do), it will be possible to make semantic mistakes which can lead us to think that coherent and complete explanations have been provided when, in fact, they have not. (I intend to return to this question in Part 2.)

matter of degree. Let SoR be the referential unit of measurement. As with the strength of concept-to-concept links, I will measure the strength of concept-to-object links on a scale from 1 to 4.

SoR=1 when:

- in background theory terms, an object is identified as *definitely* an apple;
- in cognitive-psychological terms, a concept (image and wordform) in STM and one in LTM are judged an *unquestionable* match.

In philosophical terms, the word "apple" is used to pick out an instance of the natural kind /Apple/. In lexical semantic terms, since there are no lexical pairs involved (unless in a very uninfluential role), there is no phenomenon to study (a position which naturally follows from the Saussurean structuralist view of language which separates meaning from reference, and ignores reference). In neurophysiological terms, (an account that will be developed in Part 2), the neural image of the wordform "apple" is associated with an amalgamated sensory engram of those sensory inputs which previously have been called apples (i.e. of the memory of the concrete concept [Apple]), and this engram is closely-enough matched by a current perceptual image to elicit assent to the ascription "apple" to the referent.^{21, 22}

²¹ This description is, for all intents and purposes, identical to Quine's description (in *Word and Object*) of a field linguist querying "Gavagai?" of a native speaker of a previously unencountered language. We, however, unlike Quine, are self-consciously trying to isolate reference from meaning whereas Quine was on the hunt for a concept of "stimulus analyticity" with which he could provide a scientifically-grounded account of meaning. Quine gave up, and became distracted by Davidson's theory of translation. We, however, don't need to give up because we understand that meaning is a matter of concept-to-concept connections, whereas image-to-concept connections are a separate matter.

²² I will leave further neurological descriptions to Part 2. I have included brief neurophysiological sketches such as this one here to give a sense of how relatively uncontroversial it is to claim that psychological phenomena have neurophysiological correlates. In Part 2, I will also have some comments on whether these correlations are best thought of as psycho-physical interactions, as the supervenience of the mental on the physical, as the theoretical reduction of psychological talk to neurophysiological talk, or as eliminativism, i.e. the recommendation to regard psychological talk as simply entrenched myth.

In these SoR=1 cases, the perceptual image component of the concrete concept, in STM, is *centered* in the engram image component, in LTM, of that same concept, the two images associated with the same wordform (which is also determined by an image to image LTM/STM match, these images being images of the wordform itself).

At the neural level, nothing is going on but the graded match of sensory images.

SoR=2 when:

- in background theory terms, an object is identified as *probably* an apple;
- in cognitive-psychological terms, a concept in STM and one in LTM are judged a *likely* match.

In these SoR=2 cases, the perceptual image component of the concrete concept, in STM, is located *somewhere between the center and the periphery* of the engram's image component, in LTM, with matching wordform.²³

SoR=3 when:

- in background theory terms, an object is identified as *possibly* an apple;
- in cognitive-psychological terms, a concept in STM and one in LTM are judged an *approximate* match.

In these SoR=3 cases, the perceptual image component of the concrete concept, in STM, is located *somewhere close to the periphery* of the engram's image component, in LTM, with matching wordform.

²³ For a more developed use of the metaphor of center and periphery, its association with physiological psychology, and its application to several important issues in both analytic philosophy and in linguistics, see Peter Gardenfors' work on conceptual spaces (2000, 2014).

SoR=4 when:

- in background theory terms, an object is identified as something *resembling* an apple;
- in cognitive-psychological terms, a concept in STM and one in LTM are judged *nearly* a match.

In these SoR=4 cases, the perceptual image component of the concrete concept, in STM, is located close to but outside the periphery of the engram 's image component, in LTM, with matching wordform.

Let a *referential evaluation unit* be for an object (i.e. fully instantiated) statement, and have this form:

<P, O, C(*o,w*), SoR>, <Person, Occasion, (concrete) Concept, object-to-word strength of reference>.

The questions to ask, whether in the single-match or selection-from-a-list match, are pretty obvious. The questions are, and must be, both asked and answered in terms of background theory, and so the keywords to distinguish the four categories of SoR are, in sequence, definitely/unquestionably, probably/likely, possibly/approximately and resembling/nearly.

The same set of questions could be asked in experiments in which a word (rather than an image) is presented first, and the task is to match an image to the word. Again, one variant form of this experiment would be an unprompted experiment, e.g. in which the subject was asked to draw an image of an object corresponding to the word. And another variant would be a prompted experiment in which the subject was asked to select an image (or images) corresponding to a word, from a set of images. Since the word-to-

object (wordform-to-image) association is bi-directional, these two sets of experiments should produce approximately the same results. However, persistent differences may appear in the association of a specific wordform-to-image pair, and these differences may mark cognitive issues specific to individual subjects, or differences that have made their way into the broader linguistic community.

Abstract Concepts as Concrete

Orthogonal to the concrete/abstract distinction, for ordered concept pairs, is a distinction between immediate and discursive awareness of the meaning set of the first concept. As we learn additional inferences and generalizations for a new concept, i.e. new pairings of additional concepts with that first concept, we add them to our meaning and association sets for that concept.

In using that new concept, at first, we may have to consciously rehearse the new inferences and generalizations. But eventually, we no longer have to do this. Without bringing these conceptual relationships into consciousness, we become able to use the new concept in a manner consistent with them. The second concepts in those ordered concept pairs exert their influence without our becoming aware of them.

This is a ubiquitous phenomenon. Without it, every statement we make would require us, of both its subject and predicate concepts, to "run through" the inferences and generalizations associated with it, to make sure that there are no immediate or transitive inconsistencies created by the statement's combination of its own pair of concepts.

I will call this immediacy of awareness of a concept's set of relationships to other concepts a *concept's extended gestalt*. The referential use of concrete concepts has its own gestaltist immediacy built into the word-to-object association. We intuitively

recognize this fact. But I am here pointing out that concept-to-concept associations also acquire their own gestaltist immediacy. We don't have to "figure out" that dogs are animals, or that integers are numbers, or that thoughts are mental. We know that they are. Indeed, the best account I can give of what knowledge (or belief) is, is that it is the set of inferences and generalizations for a given concept that we are immediately aware of, and aware of without needing to be consciously aware of the second concept in those concept pairs.

On the immediate-to-discursive "figuring it out" continuum, abstract concepts with immediacy are *conceptual objects*. Their meaning sets have become reference sets as well, means of immediately picking out these concepts as referents of specific kinds. These abstract concepts have become abstract objects as well.

This is the origin of our sense that many non-physical concepts are kinds of non-physical things that really do exist, just not in space and time. Many mathematicians believe that mathematical objects exist, and that when a new mathematical object is introduced into discourse, it is because the object has been discovered. Other mathematicians do not share this belief. For them, mathematical objects are constructed, not discovered, and they will refuse to incorporate into their universe of discourse mathematical objects which are not constructed according to rules of construction they restrict themselves to. For the first kind of mathematician, for example, the infinity of prime numbers really exists; for the second kind, it is just a term used to indicate that a construction process (i.e. adding one until an integer divisible only by itself and one is reached) has no end point.

Of course, purely abstract or purely concrete concepts either do not exist, or are extremely rare. Concrete concepts are ones whose sensory images (images of other than wordforms) will readily prompt assent to an ascription of their corresponding wordforms.

Abstract concepts are ones whose connection with sensory input requires transitive connections across a series of increasingly concrete concepts. However, in the case of abstract concrete concepts, the immediacy of their recognition as kinds, and of abstract objects as instances of those kinds, coalesces meaning and reference, and frees us from the need to process those concepts in a plodding, discursive way.

Across persons and time, the inferences and generalizations that are part of the gestalt of a concept will vary. To the degree that overlap exists, easy communication is possible. To the degree that it does not, we can invoke a step-by-step discursive process in which one person either agrees to a chain of inferential connections, or the two parties agree to disagree about them.

Clearly, timed-response measurements of assent (or dissent) to the co-ascription of a pair of concepts can be made in a controlled-experiment situation. Students in a new field will be slower; they will have to "figure out" the new concepts as they go along. Experts will be faster; they will just "see", or just immediately know, what those inferential connections are, out along multiple links in a chain of transitive connectivity.

Describing the process in which patterns become particular things, I would say that it is the process in which reference to kinds of objects is established as meaning inferences crystallize in an Aha! moment. It's not just that an object is recognized; it's that an object is recognized as an instance of a kind. A special case is the recognition of a new kind, of the first instance of the new kind which, like all kinds, does not exist in conceptual isolation.

So far, I have associated matters of fact with statements in which referential constraints dominate inferential ones. But we can now see that -- at least for the ontological realists among us -- there can be matters of abstract fact, or of mental fact, not just matters of

physical fact. The frequent use of a concept within a context of discourse will, over time, create an extended inferential and associational gestalt for that concept. Over even longer time, the use of that concept will feel, subjectively, as if it were prompted by awareness of something, of some thing. It is no longer just a concept which one can use with the facility of immediacy. It is a concept whose instantiations seem to be real things, whether or not they are real physical things.

Measuring Semantics

The interplay of quantified and fully-instantiated statements, in universally and existentially-quantified statements, and in purported counterexamples, is the way to see meaning and reference interacting to produce semantically valid statements, statements about concepts and statements about objects, statements about universals and statements about particulars. This is the dynamic process which constitutes *semantics*, the cognitive awareness of things, including linguistic things, as instantiations of a set of categories constrained, mutually, by lexicalized engrams and by ordered pairs of concepts.

The more concrete our concepts are, the farther removed they are from being influenced by the inferences in their meaning sets, and so the more their combinations in statements are subject only to the constraint of their lexicalized image's conformity to the lexical engram for those concepts. The more, we will be inclined to say, the statement a pair of concepts appears in is true simply because we can look and see that the statement correctly describes its referent.

But as our concepts become increasingly abstract, further removed from direct sensory experience, the more their truth or falsity is also constrained by inferences, especially strong inferences, in their meaning sets, and so the more we are inclined to say that they are true (when true) because that's just what those words mean, and that any deviation is

evidence of semantic error.

In the language of background theory, some statements will be taken to be simple descriptions of what is right in front of us. Such statements are said to be true because they report on what is really the case, on what the referent of the statement is really like. They are said to be true as a matter of fact. Concrete concepts like these, then, are sometimes able to force revision of inferential components in their own meaning sets and in the meaning sets of related concepts, by enabling counterexamples to those universal quantifications. This is how reference influences meaning.

To determine the meaning of a concept, we need to evaluate ordered pairs in which the concept is the first one. Consideration of counterexamples is a way of evaluating them. Each universally-quantified statement, for a given referent concept, is an inference. If we accept a counterexample, we have rejected the inference. If we reject a counterexample, it is for one or two reasons. Either we believe that the indicated object is not of the indicated kind, or that the indicated property is not of the indicated feature, or else we believe that a pair of inferences, obtained by the transitivity of ordered pairs, contradict one another.

Checking for inconsistencies is another way of evaluating them. Via the transitivity of these ordered pair relationships, the implications of a statement's co-ascription of a pair of concepts to the same object, at the same time, can be made explicit. Back-and-forth arguments concerning such putative inconsistencies are the other side (along with arguments about reference) of the always ongoing process of communicating with others in one's own linguistic communities. And, as I have repeatedly emphasized, these two sets of constraints on our linguistic behavior is what constitutes the semantics of language and the relationship of language to the things it is about.

Conclusion

The MRRT holds.

Meaning is distinct from reference, but not isolated from it. Meaning relationships eliminate patterns of referential behavior from our speech, patterns of counter-example co-ascriptions which would otherwise be acceptable. We could not successfully refer without these constraints. Without them, "anything goes" would be the rule for reference to particulars -- physical, mental or abstract.²⁴

And reference is distinct from meaning, but not isolated from it. Lexicalized images can force changes in the strength of various inferences in the meaning sets of those and other lexicalized images, thus altering the meaning of those concepts. Without them, "anything goes" (short of inconsistencies) would be the rule for the ordered concept pairs in inferences.

When reference sets dominate, statements are said to represent matters of fact. When meaning sets dominate, matters of meaning. Changes to dominant reference sets can force changes in meaning sets, i.e. to the set membership criteria for those concepts.

²⁴ There is another way of looking at this, however. Some philosophers and linguists would say that, without those meaning constraints, our referential use of concepts would still be constrained. It would be constrained by the "natural kinds" that really exist in the world our language is about. (See (Brzović, 2018), (Bird & Tobin, 2018), (Smith, 2015), (Brown, 2016), (Johnston, 2016)). These natural ontological categories would be instantiated by things which are real instances of those real categories. And however neurally-mediated our contact with those things might be, our brains can't conjure up any set of categories whatsoever. Nature "has joints", and it is the job of language to cut Nature at those joints.

My problem with this perspective is epistemological. If Nature has joints, the question remains of how we can find out what those joints are. Perhaps biologically innate patterns of sensory interpretation and integration -- the foundations of the work which goes on before we reach concept-mediated conscious awareness -- is that direct link to Nature. I have no problem with that, as far as it goes. But it does leave the judgment of where those joints lie dependent on the contingency of the kind of brains and peripheral nervous systems we have evolved, and the kind of language we have evolved, and these are the only tools we have for observing and thinking and talking about those joints.

Changes in dominant meaning sets can force changes in reference sets, i.e. to the range of things to which those concepts would be said to apply.

To generalize: the semantics of concepts is an interplay between set membership and set inclusion. In a semantically perfect language, there would be no conflict; the two forces of meaning (as expressed in set inclusion judgments), and reference (as expressed in set membership judgments) would be perfectly coordinated. But because of the graded nature of both meaning and reference, there are no such languages. As peripheral inferences of different concepts overlap with one another, there is a dynamism in which some inferences weaken and others grow stronger, and the same for peripheral areas of the perceptual images of concrete concepts.

So lexical semantics cannot make a dichotomous distinction between matters of fact and matters of meaning, between statements belonging in an encyclopedia and statements belonging in a dictionary. It cannot because there is no such dichotomy.

Reference works only because meaning constrains the concepts that can be conjoined in the inferences that make up the category membership criteria for the groups of individuals that those concepts refer to. Without those constraints, reference would be impossible.

Meaning works only because reference constrains the ordered concept pairs which constitute their own membership criteria. These referential constraints make some putative concept-pair-described scenarios conceivable and others not conceivable. Without those constraints, meaning would be impossible.

Together, meaning and reference constitute the semantic web of concepts, on the basis of which syntax weaves together basic statements into more complex sentential forms, thus

making it possible for us to communicate with one another without the tedious need to individually list basic statements and their truth-functional relationships.

As a final observation, TM integrates the perspectives of truth-functional and structuralist semantics.

The *reference sets* of concrete concepts, including abstract concepts which have acquired the immediacy of extended gestalts, account for our background theory belief that language is *about* something, and that what statements are about contributes to the determination of their truth or falsity.

The *meaning sets* of all concepts are expressed in inferences which, as ordered pairs of concepts, constrain the co-ascription of the pairs of concepts that we produce in our statements, and knit our concepts into structuralist webs of co-ascriptional constraints, lifting language beyond the level of a collection of discrete designation symbols.

The *association sets* of unordered concept pairs express relationships among them whose reference sets lead us to co-ascribe them, but which are at too high or too low a level of abstraction to be formulable as (universally-quantified) inferences. But in living every day, we need existential generalizations as much as, but for different reasons than, universal generalizations.

And sometimes we need to assign values to all our variables, and *fully instantiate* the conceptual patterns we become aware of. Sometimes we need to talk about things.

This essay will be followed by a Part 2 in which a neurophysiological foundation for this semantic theory will be proposed, and will be compared to and contrasted with several existing neuroscience of language semantic theories.

Bibliography

Note: articles from the Internet Encyclopedia of Philosophy, and the Stanford Encyclopedia of Philosophy, are survey articles. Note also that all dates for articles from the Internet Encyclopedia of Philosophy are the dates I downloaded the articles. There are no dates provided by the IEP itself.

- Bickerton, Derek. *More than Nature Needs*. Harvard University Press, 2014.
- Bickerton, Derek. *Adam's Tongue*. Hill and Wang, 2009.
- Bird, Alexander & Tobin, Emma, "Natural Kinds". The Stanford Encyclopedia of Philosophy (Spring 2018 Edition).
- Brown, Curtis. "Narrow Mental Content". The Stanford Encyclopedia of Philosophy, (Summer, 2016 Edition).
- Brzović, Zdenka. "Natural Kinds". Internet Encyclopedia of Philosophy. Downloaded 2018.
- Burge, Tyler. *The Origins of Objectivity*. Oxford University Press, 2010.
- Chierchia, Gennaro & McConnell-Ginet, Sally. *Meaning and Grammar*. MIT Press, 1990.
- Dehaene, Stanislas. *Consciousness and the Brain*. Penguin Books, 2014.
- Eco, Umberto & Santambrogio, Marco & Violo, Patrizia (eds). *Meaning and Mental Representations*. Indiana University Press, 1988.
- Gardenfors, Peter. *Conceptual Spaces*. MIT Press, 2000.
- Gardenfors, Peter. *The Geometry of Meaning*. MIT Press, 2014.
- Gove, Philip B. "Linguistic Advances and Lexicography". *Word Study* (October, 1961), pp. 3-8.
- Harris, Randy Allen. *The Linguistics Wars*. Oxford University Press, 1993.
- Hickok, G. "The Functional Neuroanatomy of Language". *Physics of Life Reviews*, 6, (2014), pp.121-143.
- Hickok, G. & Poeppel, D. "The Cortical Organization of Speech Processing." *Nature Reviews Neuroscience*, 8, (2007), pp. 393-402.
- Hoffman, P. & Lambon Ralph, M.A. "Reverse Concreteness Effects Are Not a Typical Feature of Semantic Dementia. Evidence for the Hub-and-Spoke Model of Conceptual Representation". *Cerebral Cortex*, 24, (2011), pp. 1767-1777.
- Hughes, G. & Cresswell, M. J. *An Introduction to Modal Logic*. Methuen, 1969.
- Jackendoff, Ray. *Semantics and Cognition*. MIT Press, 1983.
- Jackendoff, Ray. "Conceptual Semantics". In [Eco, et al. 1988], pp. 81-98.
- Johnston, Thomas. *Analyticity, Meaning, Is and Ought: the Intensional Status of Evaluative Sentences*. Doctoral dissertation, 1979.
- Johnston, Tom. *Bitemporal Data: Theory and Practice*. Morgan-Kaufmann, 2014.
- Johnston, Tom. *Transcriptions of Philosophy Notes: 1972 – 1979*. tm1972.com, 2015.
- Johnston, Tom. "Externalism and Internalism 2: More on the TM Perspective". tm1972.com, 2016.
- Kemmerer, David. *Cognitive Neuroscience of Language*. Psychology Press, 2015.

Kripke, Saul. *Naming and Necessity*. Blackwell, 1980.

Lakoff, George. "Cognitive Semantics". In [Eco, et al. 1988], pp. 119-154.

Lau, Joe & Deutsch, Max. "Externalism About Mental Content". The Stanford Encyclopedia of Philosophy (Winter 2016 Edition).

Lenci, Alessandro. "Distributional Semantics in Linguistic and Cognitive Research". *Rivista di Linguistica* 20.1 (2008), pp. 1-31.

Lenci, Alessandro. "Will Distributional Semantics Ever Become Semantic?" 7th International Global Wordnet Conference, January 28, 2014 (Tartu).

Levelt, W. J. M. "Spoken Word Production. A Theory of Lexical Access". *Proceedings of the National Academy of Sciences*, 98, (2001), pp. 13464-13471.

Loux, Michael. "Particulars and Their Individuation". In [Loux, ed, 1970], pp. 189-203.

Loux, Michael. "The Problem of Universals". In [Loux, ed, 1970], pp. 3-15.

Loux, Michael, (ed). *Universals and Particulars*. Doubleday & Company, 1970.

Lyons, John. *Linguistic Semantics*. Cambridge University Press, 1995.

MacLeod, Mary C. & Rubenstein, Eric M. "Universals". *Internet Encyclopedia of Philosophy*. Downloaded 2018.

McKay, Thomas & Nelson, Michael, "Propositional Attitude Reports". The Stanford Encyclopedia of Philosophy (Spring 2014 Edition).

Nowell-Smith, P. H. *Ethics*. Pelican, 1954.

Paivio, Allan. *Mind and Its Evolution: A Dual-Coding Theoretical Approach*. Lawrence Erlbaum, 2007.

Pulvermuller, Friedmann. *The Neuroscience of Language: on Brain Circuits of Words and Serial Order*. Cambridge University Press, 2002.

Putnam, Hillary. "The Meaning of 'Meaning'". In [Putnam. 1975.]

Putnam, Hillary. *Mind, Language and Reality*. Cambridge University Press, 1975.

Quine, W. V. "Two Dogmas of Empiricism". In [Quine, 1961], pp. 20-46.

Quine, W. V. *From a Logical Point of View*. Harper Torchbooks, 1961.

Quine, W. V. *Word and Object*. MIT Press, 1960.

Quine, W. V. *Philosophy of Logic*. Prentice-Hall, 1970.

Rosch, Eleanor. "Cognitive Representation of Semantic Categories". *Journal of Experimental Psychology: General*, 104 (1975), pp. 192-233.

Rey, Georges. "The Analytic/Synthetic Distinction". The Stanford Encyclopedia of Philosophy (Fall 2018 Edition).

Santambroglio, Marco & Violi, Patrizia. "Introduction to *Meaning and Mental Representation*". In [Eco, et. al. 1988], pp. 3-22.

Schiffer, Steven. *Remnants of Meaning*. MIT Press, 1989.

Schwitzgebel, Eric. "Belief". The Stanford Encyclopedia of Philosophy (Summer 2015 Edition).

Smith, Basil. "Internalism and Externalism in the Philosophy of Mind and Language". *Internet Encyclopedia of Philosophy*. Downloaded 2015.

Speaks, Jeff. "Theories of Meaning". The Stanford Encyclopedia of Philosophy (Fall, 2014 Edition).

van Benthem, Johan. "Modal Logic: A Contemporary View". Internet Encyclopedia of Philosophy. Downloaded 2018.