Ruth Millikan/*Beyond Concepts: Unicepts, Language, and Natural Information*
(Oxford, 2017)
Summary Notes

The notes below were produced by Dorit Bar-On in preparation for meetings of a reading group on Beyond Concepts and lightly edited by Ruth Millikan. The notes were prepared for online publication with the help of Drew Johnson. (Bullets in italics are reader glosses, comments, or questions.)

Part I: Unicepts

Introduction to Part I

0.1 Overview (p.1-2)

- Kant's q: How is knowledge possible? (p.3)

  ⇒ How is it possible for creatures with cognitive systems like ours to know the natural world?

  ● *What must the natural world be like so that creatures like us could learn/acquire information about it?*
  ● *What must our cognitive system be like so that it can represent such a world?*
  ● *How can language be informative?*

- Priority claim: natural ontology comes before cognition

(*'In the beginning there was the natural world'.*)

> "Knowledge of the world is made possible in large part by the dense clumpings together of entities with common properties that constitute, first and foremost, the endurance of individuals over time, then the existence of real though rough kinds, and finally the existence of real categories of real kinds that can support meta-induction." (p.6)

0.2 Selection Processes (p.4-6)

- Backdrop: Using contemporary evolutionary theory and its analogues in learning and cultural selection to account for intentionality.
- Background principle:

  "[I]ntentional explanation in psychology, explanation by reference to beliefs, desires, and so forth, is explanation by reference to a kind of engineering principle exemplified in mechanisms that have been formed by selection processes, natural selection, learning, cultural selection. By this principle, an organism's thought or action is usefully governed, in part, by internal stand-ins exemplifying structure-preserving mappings of affairs in the environment – 'a functioning homomorphism, which is a two-word definition of a representation' (Gallistel and King 2009: 55)." (p.4)
- The accuracy of an intentional representation is a matter of whether it corresponds as required for the functions of mechanisms designed to use the representations in producing and guiding inner and/or outer behavior.

  ⇒ An important kind of human representations are stored without immediate action or specific practical use – let alone any crudely "biological" function – in

sight. Specialized selectionist principles are at work during perceptual and cognitive development.

- On proper function:
  "[P]roper functions are effects of devices that have, to speak strictly, been *retained* (not designed) despite selection pressures and that continue to be duplicated or reproduced because they are producing these effects." And "we can also consider as a proper function any effects that a reproduced device has that are required for Normal functioning of a cooperating device" (p.6)

  ex.: A proper function of an anhinga's wings is to spread out to form shadows that make minnows more visible; the wings were not designed for this, but one of the anhinga's functional behaviors is designed to use this effect.

0.3 Ontology and Language (p.6-7)

- (See quotation above.) (p.6).
- "The manageable natural world as first confronted by creatures such as us consists largely in irregularly shaped clusters in property space of entities having correlated boundaries, perhaps fading into other clumps rather than clearly separated."
- (*So there are joints in nature but not clean ones, and correspondingly, no clean representational carving of nature at its joints in thought.*)

- Misleading comparison with the 'digital' character of language: makes it "look as though the first job of thought was "to classify, to section the world as the surface of language is sectioned".
- *So, if thought is homomorphic with the natural world conceived as above, then **it** will partake in the 'messy' character of the world it represents. This is necessary, if thought is to play its role in learning and the acquisition of knowledge.*

- "One discovery from these explorations is the importance of direct reference and how to get on without an analytic/synthetic distinction" (p.7).

0.4 Unicepts and Unitrackers (p.7-9)

- Thought's primary job is
  "not that of artificially classifying but of *locating* real properties and real entity clusters in the distal world, then in learning to reidentify or 'same track' them when encountered again in experience so as to find and follow fruitful paths for induction." (p.7).
  "The animal's first job is to keep whatever part of its distal world it would learn about *in focus*, recognizing what is the same as the same and what is different as different. Otherwise, nothing can be learned about that world over time nor can any learning be applied." (p.7).
- This job of 'sametracking' is facilitated by having *unitrackers* and *unicepts*.

Unitrackers: Internal mechanisms whose function is to enable an organism to track real patterns of sameness/similarity in the environment across space and time so it can learn about them inductively.

Unicepts: Elements of representations, which are, in turn, 'internal stand-ins homomorphic with environmental states of affairs'.
- Unitrackers are cognitive faculties enabling an organism to reidentify what is the same again thus enabling inductive learning/reasoning. Unicepts are elements of cognitive representations. The traditional notion of "concept" conflated unitrackers (faculties for same-tracking) with unicepts (elements internal to judgment and belief).
- Unitrackers and unicepts are not shared across individuals, though their *targets* – what they are used to track/represent – are.

Ch. 1 A Clumpy World

1.1 Overview

1.2 Real Kinds (p.11-15)

- "Not just any possible world could be thought about or talked about. A known or thought-about world must be organized in a suitable way" (p.11).
- Chapter 1 aims to outline "a certain kind of organization that the world clearly exhibits", describe its importance to a 'cognizing animal', and partially explain "how the world becomes organized in this way.
- Starting point: an imagined multi-dimensional graph representing a logical space with each dimension corresponding to a different 'property dimension', and with dots representing 'various ordinary physical objects in the world' at a given time. (p. 12).

Note:
(1) When every object has been represented, "all but very small areas on the graph will remain empty".
(2) Non-empty areas will contain mostly "clumps of dots that were in close proximity along multiple dimensions" (e.g., the rabbit clump, the Gothic cathedrals clump, CD disks, etc.).
- "The world of physical objects is to a large extent filled with clusters each having densely interlocked properties, clusters that are for the most part distinctly though not always perfectly separated from one another. This kind of structure is what underlies the success of everyday induction …" (p.12ff).
- Clusters fail to "fall under necessary descriptions", but they typically "exhibit properties that are "both common to many members and well distanced from" those found in other clusters.
- The availability of "different methods of identification for recognizing clusters is "indispensable to the accumulation and use of knowledge" concerning them.
- "Real kinds": clusters that are "formed within any general category of things, within any multidimensional logical subspace", where the clustering is non-accidental – has a univocal reason.
- Though it is objects at times that we encounter as we gather information, what we are gathering information about are "the individual objects of which they are stages or the kinds of object of which they are stages".

1.3 Reproduction and Mass Production (p.15-17)

- Much of the clustering can be explained through processes of reproduction and mass production.
- The reason only a limited variety of things exists is that the 'stabilizing processes' that take over once some things 'happened to find their way into our world', built in such a way so as to persist, use up "materials and locations … removing these materials from availability for the constitution of other things".

1.4 Historical Kinds (p.17-18)

- 'Clots' are kinds whose members are 'glued together' in that they "have their typifying properties in common owing to causal connections among them, coupled with support from a common persistent environment".
- These are 'historical kinds' in that, like individuals, they 'bundle sets of properties together repeatedly "not merely owing to causal dependencies among these properties but owing to a common historical origin or source". They can include artifacts such as screwdrivers or hard drives.

1.5 Individuals (p.18-21)

- Individuals are historical entities – clots composed of time stages of that individual.
- The distinction between historical individuals and historical kinds is slippery, but what is of interest for understanding cognition is their striking similarity. Both exhibit recurrent clustering of properties – stable over time and space – that "create opportunities for inductive knowledge".
- Linguistically, we see terms that behave like *names* of historical kinds. (E.g. "Dogs bark"/ "The/a dog barks". "Sentences of this kind *seem to express thoughts about an unindividuated property clump,* thought of merely as something that can be encountered here and there or elsewhere and that will or may show these and those traits when encountered.*"* (p.20)
  - *(See earlier q re relation between language, thought, and world.)*
- There are also 'social individuals' (insect colonies, lion prides, clubs, societies, universities …), all of which can be kept track of.
- "The existence of historical individuals and kinds is the most fundamental source of usable natural information, which lies at the foundation of all cognition and all language use. [They lie] … at the center of cognitive life. They are the main structures that make thought and language possible."
  - ⇒ ***Radical semantic externalism:*** It's the worldly kinds that give structure to 'meanings'. Meanings do not 'carve up possible worlds but are 'structured and maintained almost entirely by the actual world'.

1.6 Eternal Kinds (p.21-22)

- Eternal kinds are ones whose members cluster together for non-historical reasons. "[T]he possibility that more members of the kind may appear in history at separated places and times remains a permanent possibility".
- Examples: "samples of organic matter like water and quartz and glucose" (p.21).

1.7 Shapes and Divisions of Historical Kind Clumps (p.22-24)

- Historical kinds are not held together merely by common properties nor are they governed by exceptionless laws. They "evolve as they move forward in time, and, unlike most individuals, they may branch." (p.23)
- The natural world is 'messy,' with clumps of clumps that come in a variety of shapes with no neat hierarchy.
- This is the world "with which language and thought have to cope". But in fact language speakers pick up names for individuals and for real kinds very quickly (*what's known as 'fast mapping'*). And the cognitive trick involved arises 'well prior to the emergence of language' and even to human cognition. Nonhuman animals have a great variety of reliable ways of recognizing individuals and real kinds.
- That worldly kinds admit of being thus recognized is essential to the possibility of both cognition and language.

1.8 Real Categories (p.24-26)

- 'Real categories' are kinds of kinds. Real kinds "correspond to correlations among determinates that are found together, each characterizing many or most members of the kind. Real *categories* correspond to correlations among determinables that are found together, *just one determinate from each* characterizing most members of the category". (p.24)
- "The members of a real kind have various determinates in common; the members of a real category have various determinables in common, each having a stable value for each category member." (p.25)
- The distinction between real kinds and real categories is relative.
- *See examples p. 25f.*

Ch. 2 Direct Reference for Extensional Terms

2.1 Overview

2.2 Conventions of Language (p.27-30)

- In one *simple* sense, a convention is just a pattern of behavior that has been handed down from person to person and that, insofar as it has a function, that function could be served equally well by other patterns. (=*arbitrariness*).
- Convention in this sense (unlike Lewis') doesn't entail that the pattern is (even nearly) universally followed, or that it always solves a coordination problem (though it can), or that it's always mandatory/obligatory.
- Contra Lewis, simple conventions that do solve coordination problems can be learned from experience; passing them on does not require Gricean inference.
- Linguistic conventions are simple conventions of this kind.
  "The repeated, handed-down, problem-solving patterns that compose the core of a conventional language in use involve not merely *acts* of the speaker but *cooperative* acts of the speaker, and not merely acts of the hearer but cooperative acts of the hearer." (p.29)
- The "help" example (p.29): The patterns of behavior passed down include both speaker act and hearer response. "All of these parts of the pattern must be in place enough of the time if the convention is to survive".

- What is of interest here is Normal use of language, “how language works when it is alive and reproducing itself”.

2.3 Following Precedent (p.30-32))

- A good start on the semantic ‘poverty of stimulus’ problem in following precedent, as posed by Lewis, can come from “understanding how the clumping and peaking features of the world itself serve to constrain as well as to support cognition”.
- Two implications: (1) “there are no such things as determinate ‘rules of language’; (2) “the truth rules for a language cannot be fully determinate”. Especially in the face of worldly changes in boundaries between clumps, “precedent may no longer be able to determine correct usage”.

2.4 Direct Reference to Clumps (p.32-34)

- Against description theories: clearly we identify individuals through a variety of characteristic properties, not all shared between users of a name.
- Real kinds make themselves easily available to cognition, as they are recognizable in a variety of different ways.
- Given that there are many ways to recognize real kinds, there will not be only one set of descriptions for the application of a name.
- “As a user of the name ‘dog’, all that is required that you have some relatively reliable ways of recognizing instances of the dog cluster” (p.33).
- What makes our thoughts/words be about the same thing is just that it is the same thing our thoughts/words identify with fair reliability.
- A clump is a scattered individual spread out in time and space. The proximity of each to the next is not spatial, of course, but multi-property similarity. The arguments that names for individuals refer to them directly rather than referring through descriptions apply as well to real kinds, to clumps. Fido “is a dog” if he is somewhere in the dog clump. The word “dog” points to but does not describe that clump. Different people may recognize that this or that must be in the dog clump in different ways and in different ways at different times. Thinking of an individual or of a kind does not require knowing anything in particular about it.

2.5 Identifying Through Language (p.34-35)

- All identification is recognition through signs, signs of signs, and so forth. This is so for recognition through perception, through language, and through natural signs.
- Names for individuals and real kinds are directly referential; merely having a word can be enough to have a thought of its referent, to begin collecting info about the referents.
- One just needs to be able to recognize different tokens of the same name *as such* - to know that the tokens are about the same thing.

2.6 Real Definitions (p.35-36)

- Concerning real kinds; there is a reason why members of a real kind are similar to each other.
- Real definitions of kinds concern the explanation of why members of the kind are like each other.

- "They explain the nature of what, operating in the background, is holding the extension of a term together" (p. 35).
- But you do not need to know the real definition in order to use a term to talk about a thing.

2.7 Names for Properties (p.36-37)

- Like with real kinds, the structure of reality provides an anchor for the names of properties.
- A property space contains natural anchor points. Some of these are 'peaks', arrived at when approaching from any of various property space directions ("pure", "absolute", etc.).
- Some names are tied to directions rather than peaks. ("redder", "flatter", etc.). "Big", "tall" → anchored to directions; no boundaries. "small", "short" → anchored to direction; bounded.
- Many anchoring points are eternal (not hitched to particular location): "perfectly circular", "pure water", etc. Other anchoring points are 'historical', fixed by stable features of historical kinds' "pear-shaped", "furry", "bony", etc. Basic words for colors are anchored in human perceptual system. Some adjectives stand for property complexes ontologically similar to real kinds, e.g. "Rot", "mold", "anger", "bankruptcy" - these are common syndromes of properties.

2.8 Boundaries and Slippage (37-39)

- Property names anchored by peaks do not have boundaries. They describe whatever is close enough to their anchoring points for present communicative purposes. This is a matter of context. This dissolves the Sorites paradox.
- Names for real kinds are anchored to clumps, bumps, ridges, etc. But, as clumps may be elongated, branch, be parts of other clumps, it is not always determinate which stages/branches are included under a name.
- Names for things will have edges (application rules) only as sharp as the things themselves. "If there were no natural boundaries or clear gradient shifts in the space-time actual world there would be nothing whatever to separate the extensions of names from one another and neither communication nor thought would be possible at all" (p.38).
- Word use changes over time and context. There is only a fuzzy distinction between the 'literal' lexically encoded use of a word and a 'widened' or 'narrow' use.
- "the surface of language is digital; the world is not". But: "In many dimensions in many places the world proves close enough to digital to tame language, though not to control it strictly" (p.39).
- The world changes over time, so that boundaries that were once pretty clear may become unclear. One natural result is changes in the meanings of many words and other expressions over time

2.9 Communication with Names for Clumps and Peaks (39-41)

- Nothing beyond pure reference is automatically carried from speaker to hearer for extensional terms. So how do words communicate?

- Uses of names are anaphoric on an independent understanding of the things referred to. (But of course, the understandings are not the same).
- "There is no such thing as mere linguistic understanding. Understanding language is *buried* in understanding the world as a whole" (p.40).

- *Question re p. 40 "any characterization that it is a speaker's purpose to impart with a word must be reconstructed by the hearer from his own knowledge and from the context": doesn't that depend on **having** richer content in mind to impart/reconstruct? If 'understanding the world' is itself like the understanding of a bare 'this' and 'that', what are we to make of even non-linguistic understanding?*

Ch. 3: Introducing Unitrackers and Unicepts

3.1 Overview

3.2 Initial Examples of Unitracker Function (p.43-46)

- All higher species must somehow be able to use representations of their environment to guide their behaviors. Unitrackers are posited as mechanisms that help to perform this function.
- "A unitracker for a thing takes in a diversity of proximal stimulations over time and interprets or translates them as signs carrying information about one and the same thing" (p.43).
- A unitracker 'funnels' factual or affording info about the same either into immediate use, or into storage by means of new connections of its proprietary unicept, thus bringing together different bits of info about the same thing for use in inference or action.
- "A unicept is a structure that forms a stable link, originated by its partner unitracker, between items of knowledge that are about the same . . . The referent of a unicept is its unitracker's target, what the unitracker has been designed to track." (p.44). Example: consider all the ways you can recognize your friend. Each corresponds to a different procedure for same-tracking your friend.
- *Fn. 1*: In earlier work, the notion of "empirical concepts" was used. It has now been replaced by two notions: *unitrackers* – mechanisms with the function of translating information-bearing sensory patterns into connections between unicepts – and *unicepts* – elements connected to other unicepts or to action dispositions to form intentional attitudes ('factic') or representations of affordances ('affording'); parts of complexes apt for storage of information or action potentials. Both are embodied faculties or systems that have functions/display abilities; they are two aspects of cognitive function.

3.3 Discarding Concepts (p.46-49)

- Names involve no description at all of the thing named.
- Understanding a name requires activating unitrackers for translating information from linguistic signs into inner representations.
- States of affairs are represented by connections among unicepts.
- Unicepts are used in storing beliefs and affording knowledge (and in its occurrent use).

- Unitrackers are used in translating sensory data into beliefs or affording knowledge.
- On factic unicepts; understanding a name does *not* require either
(a) Particular knowledge of the character of the thing named.
(b) Grasp of any particular way to identify what is named (other than by the name).
(And neither (a) nor (b) need be shared by anyone possessing a unicept for the same thing.)
- Neither unicepts nor unitrackers are shared - unique to each individual; particulars, not types.
- Unicepts are *not* concepts.
- Unicepts: Particulars, idiosyncratic, reference determined by the design history of its unitracker. Unicepts can be redundant or confused or equivocal.
- Traditional Concepts: Types, publicly shared, individuated by dispositions, can be empty but not confused, equivocal, or redundant.

3.4 Details on the Nature and Function of Unicepts (p.49-51)

- Unitracker: Input network, or mechanism for adding, changing, strengthening, or weakening connections of a unicept to other unicepts or action potentials.
- Same-tracking involves a designed capacity to recognize sameness itself of the target; not *merely* being able to track something for continued use.
- Collecting info about the same *as same*, is central for making mediate inferences.
- Unitrackers involve many unitracking methods, none of which is central for fixing the target. This is a feature, not a bug; it allows for equivocation in unicepts.
- Tweedledee/tweedledum and Cicero/Tully examples
- Same-tracking is not the same as classical identifying. No such thing as identifying *tout court*. Only identifying one sign as a sign of the same thing some other sign is a sign for.

3.5 Lifespan and Growth of Unitrackers and Unicepts. (p.51-52)

- Some unicepts/unitrackers are short-lived. (Keeping track of a glass at a party).
- Unicepts retained over time may accumulate connections to other unicepts. Unitrackers may grow as well; ways of same-tracking can be sharpened and tuned for greater reliability.

3.6 How Names Connect with Unicepts (p.52-53)

- Dulling the intuition that there is a difference between direct perception and learning information from hearing true sentences:
  - ⇒ No confronting things-in-themselves. Even perception is proximal. Using glasses, binoculars, telescopes is no different in kind from 'direct' perception. Volt meters, geiger counters, etc. - not different in kind either. Baby crying as sign of hunger, dog bark as sign of someone at the door. A friend's remark that it is raining. Words in a book on Napoleon – also not different in kind from direct perception. (This is unpacked and supported in Ch. 14)

- Names continue to be reproduced because they play a role in transmitting info about the things they name. Names are designed to be inputs for unitrackers. A repeated word shows sameness of reference/extension by its own sameness.
- A unitracker can include and is sometimes exhausted by the ability to recognize incoming information about a thing just by seeing or hearing its name. Two people can recognize information about a thing by means of the same name without having any further similarities in their respective unitrackers or unicepts for that thing.

3.7 Language in Unicept Development (p.53-54)

- Phonological structure enables infants to recognize when they are encountering the same word again.
- Words are made up of recombinations of a small pool of phonemes - so being able to same-track phonemes gives you the resources to same-track and collect info about everything named by the language!

3.8 Modeling Unicepts (p.54-55)

- What comes to mind when a name is heard will depend on what the learner knows about what is named. A lot of this will be probabilistic for real kinds.
- When you are told a name, this activates the relevant unicept, as well as probabilistic connections to other unicepts, depending on what you know.
- Models for many unicepts will require "not just that its connections to other unicepts are weighted but that there be a way of adjusting these weights for individual members (this particular dog) or stages (John as a teen) of its target according to what else is known about that member or stage" (p.55).

Ch. 4: Functions of Same-Tracking

4.1 Overview

4.2 Perceptual Constancy Mechanisms (p.56-59)

- Most *unitracking* mechanisms are fashioned from fine-tuning or combining *same-trackers* for various different properties of their targets.
- Not all same-tracking is unitracking, that is, a unitracker may have no corresponding unicept. Some sametrackers work *prior* to any unitracking. They have ways of tracking something as same, but where there is no need for collecting information about this thing, so there is no corresponding unicept. E.g., perceptual constancy mechanisms.
- Color-constancy mechanisms are same-trackers for colors. They can be employed as unitrachers by adult humans who have thoughts about the various colors. They can also be employed by beings lacking unicepts. They are then used merely to help in identifying further properties/objects/kinds by the presence of the color tracked.

p. 57: "I may, for example, identify lemons partly by their color without bringing on line my unicept for yellow". Sametracking can bypass the need for unicepts.

- Children have to sametrack phonemes in order to understand language at all, but learning to identify phonemes *consciously*, learning how to identify them for

purposes of learning and thinking about them, and so forming unicepts for them (as is needed to map phonemes onto letters) is very hard for many children.
- Edge-detectors are same-trackers but not unitrackers - we have no need to collect info about the *edges* that are tracked; they are just tracked to help us identify further things.
- "The central use of perceptual constancy mechanisms seems to be use *in context* for same-tracking of objects and kinds, not knowing about the properties themselves.

4.3 Self-relative Location Trackers (p.59-60)

- Another kind of non-uniceptual sametracking: Tracking spatial relations to oneself of objects seen or felt or heard. Discerning various sensory inputs as originating from the same location relative to oneself, hence as arising from one and the same object, may help you to identify an individual/kind. But to do this, you do not need to employ a unicept for that particular self-relative position.
- Tracking spatial relations to oneself is required, for instance, to achieve hand-eye coordination.
- This same-tracking is unusual in not being tracking over time, but instead tracking over space and sensory modalities.

4.4 Object Constancy (p.60)

- Object constancy also a non-uniceptual same-tracking.
- Studies have shown that object constancy does not depend on prior identification of qualitative properties - just spatial location and trajectory. - So maybe this kind of same-tracking helps in the development of qualitative property same-tracking.
- Used in identifying and reidentifying (keeping momentary track of) a particular object prior to any uniceptual involvement. First you track it perceptually, then perhaps you can identify it, connecting it with a unicept. Though they may on occasion be used to develop unitrackers. (Pylyshyn's FINSTs seem to be temporary unicepts that may lead to more permanent ones, and used to track just some particular object.

4.5 Same-tracking for Application of Unicept Templates (p.61-62)

- "A faculty that same-tracks a category is not just as such a unitracker" (p.61). But most real categories are also real kinds, and we can have unicepts for real kinds.
- Nevertheless, knowledge of a category provides a template for the development of unicepts for members of the category.
- This is because knowledge of a category provides knowledge of the kinds of questions that can be answered, hence sensibly asked, about each member of the category. Answering these questions about a member allows you to collect information about that individual, and may help to re-identify it later.
- Examples; category cat - you know the pattern of an individual cat won't change, so knowing the determinate pattern of Tabby's fur may let you re-identify Tabby.

4.6 Practical Stuffs and Affording Unicepts (p.62-66))

4.6.1 "Stuffs" (62-63)

- Three uses for unitrackers:

1. Affording: Identify structures in current environment as affording.
   Affordances: Possibilities for behaviors that would achieve certain ends. Contra current Gibsonians, this needn't be unmediated by inner representation.
2. Substance: In making factual judgments concerning a substance (= an individual or real kind or anything else being considered as something that has properties).
3. Attribute: In making factual judgments about the properties of various kinds of things.

- These uses are neither mutually exclusive nor necessarily concerning different kinds of target. (The same-tracker for gold that is used in discovering that my ring is gold is also used in discovering that gold is malable.)

4.6.2 Affording Unicepts (63-65)

- 'Property' and 'substance' are relative, not absolute, ontological categories. "Is bone a thing with properties, or is it a property, of things?" (p. 63).
- What philosophers call 'properties' can almost always be thought of as a 'stuff' we can collect affording or factual knowledge about.
- One knows how to recognize and deal with the slippery . . . the same way one knows how to recognize and deal with one's boss: through a grasp of their affordances.
- "Affording unicepts do not distinguish among properties, individuals, material, kinds, and so forth" (p.64).
- Pushmi-pullyu representations, PPRs: representations that are both descriptive and directive. Truth condition + fulfilment condition. Grasping a current affordance involves an inner PPR.
- Affording unicepts *do not* as such figure in *belief*s about the situation represented. (PPRs are more basic than belief). Which affordances you perceive is guided by current interests. Affordances *potentiate* actions in the current environment.

4.6.3 Affording Unicepts for Natural Continua (65-66)

- Affording unicepts for self-location needed in addition to non-uniceptual same-trackers for self-location. Cat and mouse; cat perceives mouse as a *thing to catch* and *as requiring to move that way (to catch it)* - so *two* affording unicepts are employed in one PPR.
- Affording knowledge of how to reach to, step on, walk to, etc. → unicepts that depend on a prior ability to same-track self-relative locations.
- Self-relative location-tracking same-trackers are unusual. The relations tracked and reidentified are not discrete, but continuous.

The relevant unitrackers = "unitrackers for natural continua" (These gather ways to recognize positions on a natural continuum – in this particular case, positions relative to oneself.

- An affording PPR: The positions on a continuum are mapped onto response continua that further such ends as touching, hitting, etc.

- Reveals another difference between concepts and unicepts. Concepts are tools for classification. But these unicepts reidentify positions on a continuum without classifying anything.

4.7 Factic Unicepts; Substantive and Attributive (p.66-68)

- Can't separate either an affordance or a PPR that represents the affordance from its time and place - a particular representation of an affordance would be a *different* representation if located at a different time or place. Time and place are selfsigns in a PPR. Time and place function as a subject term does in a sentence. The claim of a PPR is that here-now has certain affordances/properties.
- Thus no unicept is needed to represent a subject in the representation of an affordance.
- A belief that can be stored away, kept for use at another time, does require a subject term - must say something about something, these somethings being represented by unicepts.
- Humans can make judgments about things they are same-tracking and store the resulting beliefs away. This is done by assigning properties and relations to these things or assigning these things as properties to other things.

(a) Substance unicept
(b) Attribute unicept

- The difference between substantive and attributive unicepts concerns *function*, not the ontological status of the objects involved.
- Substance unicepts: Represent their objects *as substances* (think Aristotle), as things with determinate properties, as things that come in categories, hence supplied with a template. Substance unitrackers have the job of recognizing info about the substance as info about *that* substance.
- Attribute unicepts: Represent their referents as (one or more place) properties of substances. Must be developed along with unicepts/unitrackers for some contraries of the attribute (Even if only the property not-φ, that is, the property of being contrary to φ (67).
- Factic inner representations employ attribute unicepts →allow for 'predicate' negation. An operation on a representation to represent an indefinite contrary of its original.
- Corollary: attribute unitrackers are complemented with different-from trackers. Can perceive that something *does not* have a property.
- Whether something is represented as substance or attribute is relative, not an ontological difference.

4.8Two Closing Remarks (68-69)

(i) There is no specifiable set of determinables that a person has to recognize as being determinate for a substance in order to possess a unicept for that substance.
(ii) Nothing prevents a single same-tracking mechanism from supporting more than one of the three kinds of unicept functions. (But don't assume a unitracker designed for one unicept function can automatically be used for any other.)

Chapter 5: How Unicepts Get Their Referents

5.1 Overview

5.2 How Unicept Referents are Fixed: The Quarry (p.70-72)

- In the literature on how reference is fixed, it is assumed that "extensional terms correspond to concepts that are shared by all who use the terms correctly" (p.70). But unicepts are not shared. This chapter focuses on the referents for unicepts.
- Unicepts have input from their unitrackers, which originally control effects of sensory input on connections among unicepts. Connections among unicepts then help to determine further connections among unicepts (roughly, this is inference). Unicepts also have connections to action potentials.
- Classical descriptivist theories of concept content: Possessing a concept of a thing involved knowing certain properties are essential/characteristic/unique of that thing. The referent is whatever has these properties.
- The properties essential to a thing one has a concept of must themselves be conceptualized in order to have that concept, and so on. So a foundational theory of content determination for some concepts was needed (an input theory).
- Well-known causal and informational input theories (Fodor, Dretske, Stalnaker) face the problem of error. An application of a concept may fail to have been caused by or carry information about its referent.
- Teleosemantics offered a solution for simpler kinds of mental representations. "Looking to what a device was designed to do rather than to its actual dispositions leaves room for error in identification" (p.71).
- Teleosemantics gives an explanation of the content of the representation of a full state of affairs with a satisfaction condition, either descriptive or directive. No propositional structure.
- Content corresponds, first, to the proper function of the representation's producers which is to produce representations that map states of affairs according to given truth rules. The truth rules are rules that the systems using these representations are designed to rely on in serving whatever functions they may have.
- The referent of a unicept is determined by the proper function of its unitracker, what the unitracker was designed to track. Remember - unicepts are particulars, so there is no worry about individuation. Unicepts in different organisms may be equipped with unitrackers that have quite different ways of tracking the same target but so long as the targets are the same these unicepts will have the same referent/extension.
- Content is *not* fixed by (i) what one knows about the target, (ii) dispositions to apply a unicept, (iii) connections to other unicepts, or (iv) actual input.
- And unicept possession does not require (i) understanding of the target's nature, or (ii) perfect success in tracking.

5.3 Two General Principles Concerning Functions (p.72-74)

1. "Many Biological mechanisms exist just in case."

Eye-blinking reflex, antibodies, cognitive systems . . . are all there just in case they happen to be needed.

We store away information in case it might be useful. But info about the same can only be useful if it is grasped as being about the same, where this requires mechanisms for creating new unitrackers and unicepts.
2. "Specific biological functions are often derived from prior relational functions" (p.73).
Chameleon example. *Relational mechanism*: Change skin color to match the color of what the chameleon is on. *Specific adapted function*: here, a function of turning a particular shade of brown when on something that shade of brown.
The derivation of adapted functions from relational functions is the general principle involved in the 'priming' of new unitrackers by experience (p.73-74).

5.4 Imprinting (p.74-75)

- Imprinting is a candidate for a unicept humans and other animals are born with; maybe the only one.
- Imprinting allows newborns to same-track their parents, often by use of an identifying determinate.
- 'The baby [ducklings] imprinting mechanism has the relational function of setting up an affording unitracker for tracking whatever is related to the duckling as its mother" (p.74). The adapted specific function is to track Samantha (who is in fact the duckling's mother). It does this by causing the baby to same-track and follow the first moving thing it sees.
- So, a unique adapted affording unitracker is formed. It is skeletal in that only one/a few ways of tracking are involved, but it can grow. The unitracker's function is to track Mother, but it also has the function of tracking the first thing it sees. Since in odd cases, Mother may not be the first thing it sees, the unicept formed may be equivocal at the very start.

5.5 More General Mechanisms for Priming Unitrackers (p.75-76)

- Many species use odor to sametrack individuals.
- The system that generates specific odor-unitrackers, unlike imprinting, are not designed to prime a unitracker for a predetermined individual, but for whatever new individual comes along.
- Once created, the unitracker already has success criteria.

5.6 Some Mechanisms that Set Targets, Specifically, for Affording Unicepts (p.76-77)

- Reinforcers/rewards in an animal's native reward system can set up targets for unitrackers that track stuff associated with these rewards.
- These targets can be set in advance of reliable achievement. The goal is to find ways to reidentify these affording situations. General searching for repeated patterns of properties and attempts to complete property patterns in anticipation can be used to discover categories which then prime unicepts.

5.7 The Problem of Location-Detached Signs (p.77-79)

- So far, we've only considered directly perceptual methods of same-tracking, which typically require current presence of the thing tracked. It must be in the immediate environment. These signs typically show space-time relations of the

perceived to the perceiver along with other things about the perceived. These are *location-reflexive* or *'attached' signs* (signs carrying info to outer senses for ordinary perception). This works fine for affording concepts, which are activated in the here-now. But how do we acquire and test concepts of things we don't perceive as in certain space-time relations to ourselves.

- Attributive and substantive unicepts can be used to accumulate information about far-removed parts of the environment.
- *Location-detached signs*: Linguistic signs, maps, animal tracks, etc. Carry information about affairs without showing the relation of the perceiver to the perceived. Perhaps something in distant space and time. *Q*: How are targets established?
- Same-tracking names by phonological structure helps; names prime unicepts for what they name.
- The Q: "How do we locate things to samt-track and then learn how to track them without first located and following them in space-time relation to ourselves?" (p.79)

5.8 A Third General Principle: Proxy Functions (p.79)

- Distal vs. Proxy Functions

A rat's mechanism for reinforcement by sweets:

Distal Function: Bring in more calories.

Proximal Function: Produce behaviors that bring in sweets. This is an indirect, quick, and safe way to learn behaviors that bring in more calories.

- A proxy function can fail in its more distal function. A Normal condition for proper performance of a proxy function is that it performs its distal function.

5.9 Natural Epistemology for Substantive and Attributive Unicepts (p.80-83))

- What makes you confident that various same-tracking methods for a thing really do converge on one thing, or anything at all?
- Human cognitive systems use a proxy selection mechanism that selects, but fallibly, for non-emptiness and univocity in factic unicepts.
- The mechanism uses conformity of unitracker outputs with *laws of identity and noncontradiction*. These are regulative principles monitoring unitrackers.
- Factic representations allowing internal negation are needed to supply negative feedback for the tuning of unitrackers.
- Developing factic unicepts involves that one learn how to same-track subjects and predicates of judgments so as to agree in judgments with oneself.
- *Coherence* tests for *correspondence*. But not 'coherence' in a holistic sense; only within a small set of contrary spaces. So many independent realms of same-tracking can be discovered.

Ch. 6: Misrepresentation, Redundancy, Equivocity, Emptiness (and Swampman)

6.1 Overview

6.2 Failures of Biological Function (p.84-88)

- Unitrackers are mechanisms with a certain kind of function. But unlike hearts and other biological mechanisms, unitrackers' structure is (almost always) developed through learning and tuning in response to the environment (not genetic coding). Any mechanism with a function can fail to function normally.
- Distinguish *normal* from *Normal*: normality is a matter of statistical frequency. Normality is to do with "form or function of a kind that has been actively helping in the preservation of a real kind" (p. 85). It's an open question how normal Normal behavior is.
- Distinguish two kinds of failure to perform Normal functions: *internal* – due to defects in the relevant organ/mechanism – and *external* – due to failure of the environment to support or cooperate with its operation in the Normal way.
- P. 86: "[C]ognitive failures … are usually external failures. … [They are] more typically caused by an abnormal external environment" and "are quite common".
- When it comes to animals' common behaviors – reaching, chewing, washing, running – success requires a supportive environment, but support is usually forthcoming, since perception of the requisite environmental conditions is usually accurate. Other behaviors, however – hunting, stalking, mate-seeking – often fail, due to lack of the needed support or 'cooperation' from the environment.
This applies to unitrackers' functioning: they may fail to same-track and the mechanism responsible for priming, developing and pruning new unitrackers may also fail.

6.3 False Beliefs (p.88-89)

- A 'well-focused' unitracker with 'a clearly defined target' can fail to track correctly. It can mistake B for target A, with the result that there is false belief.
- [Cases of failures: (i) illusion (no A) (ii) misperception (mistaking B for A, as in twin cases); and similarly for belief]. Typical failures are due not "to malformed or damaged cognitive systems but to lack of environmental cooperation".

6.4 Redundant Unitrackers and Fregean Senses (p.89-91)

- Same-tracking mechanism can be ['dedicated']: you may be able to same-track distance so as to manipulate and move among objects yet be unable to tell what the distance is in inches and feet. No free transfer between recognizing objects as permitting affordances and objects as elements of facts.
- So called 'Frege cases' all trade on individuals having more than one unitracker for tracking a single object. (E.g., tracking "Hesperus" by noting location at sunrise and tracking "Phosphorus" by noting location at sunset.)
- Frege's explanation: we associated different 'modes or presentations' with different words that denote objects. These give the *senses* of names [and constitute different ways of thinking of the object].

  "[U]sually, no particular way of thinking of a thing, no particular knowledge of it, goes along with merely understanding a word for it. Most extensional terms refer

directly. … They also carry no particular ways of recognizing their targets – except, of course and importantly, via the names themselves." (p.89)

- There is no (shared) unicept – *our* unicept – that goes along with a given public name. Unicepts are individual mechanisms had by individual thinkers. Individual unicepts are associated by individual speakers with names. An individual speaker may fail to know that Hesperus is Phosphorus – as Frege observed, this is due to the speaker "having thoughts of one and the same thing in two ways," that is, mistakenly having two different unicepts for the same thing. Information that she gets from true sentences containing "Hesperus" is not stored using the same unicept as is information she gets from true sentences containing "Phosphorus." But this doesn't require postulating public 'senses'.
- P. 90: Unlike Fregean senses, ways of recognizing a target (methods for same-tracking it) "are not equivalent to either things that [one knows about it] or to ways that [one] thinks about [it]."
- Any mental representational system will be subject to Frege cases and the possibility of the individual believing what is in fact a contradiction. Contra Kripke (1988), it is no puzzle or paradox that a rational individual might believe a contradiction.
- "Read in the natural way, 'Johnny believes that Cicero was bald but that Tully was not' does not imply 'Johnny believes that Cicero was bald but that Cicero was not.' You can't directly translate a term that is standing for itself" (p.91). (see 10.7.2-3)

6.5 Equivocepts (p.91-93)

- Re the case of goslings imprinting on Lorenz' boots: what was the target of the unitrackers primed this way – the goslings' mother or the boots? – The target was equivocal. *Success in proximal function, but failure in distal function.*
- This can happen with human unicepts. "A unitracker will have as its function to track whatever it was that primed it, but will also have as it function to keep on tracking the thing it has been bringing in information about previously" (p.91). *This allows for a mismatch between proximal and distal function.*
- A unitracker that has the job of bringing in information about Tweedledee may mistakenly bring in information about Tweedledum. The unicept may hold information about two targets. It has become an *equivocept.* Examples of equivocepts from the start: 'St Patrick', 'jade'.
- 'Slow switching' cases: Having begun its life as tuned to water, Oscar's unicept has *become* an equivocept upon becoming tuned to twin water.
- No point in trying to settle the question 'At what point does a unicept become an equivocept?'; no sharp boundaries and no theoretical need to decide.
- In any event, unicepts' targets are not determined by dispositions but rather by proper functions. (Perhaps, in the end, equivocepts are not strictly speaking unicepts, since they haven't achieved the function of even *having* a definite target.)

6.6 Vacucepts (p.93-94)

- Lorenz' spurious imprinting with a laser-pointer: no real target, but only a 'moving' one. Would-be unitrackers can be set on non-existing targets. *Empty names* (spuriously priming a person's cognitive system for tracking something that isn't there) – e.g. "Santa Claus", "Volcan", "caloric". *Vacucepts.*

- Vacucepts can feature in psychological explanations. Like valid unicepts, – they involve "dispositions to make … identifications and … inferences" but mistaken ones.

6.7 How Unicepts Fit with Biosemantics (p.94-95)

- Representations in the sense of LTOBC were not mere 'intentional icons' (as described in LTOBC). They were intentional icons that contained elements whose functions involved that their semantic values [referents, extensions] be "identified." To identify the semantic value of an element α of a representation was to use the representation together with another representation containing an element β "in mental operations [such as mediate inference] the success of which depended on the values of their α's and β's] being the same". So p. 95: "LTOBC's representations .. *contain* [something like] unicepts as elements, differing in this way from simpler intentional icons. The basic job of a unitracker used in factual judgment is to aid in the making of 'representations' in exactly the sense of LTOBC. … Biosemantics is the foundation underlying the postulation of unicepts."

6.8 Swampman (p.95-96)

- Swampman is a physical duplicate of Davidson that resulted from a purely accidental confluence of molecules. Given our experience with humans, and given his physiological similarity to humans, you could use your experiences to make correct inductions about him. But their correctness would be a wild coincidence.
- Our dispositions to identify things as human do not determine the referents of our unicepts for the kind, human, nor the referent of the word "human." The referent of the unicept is "whatever has primed it and has been responsible in the past for apparent success in collecting information about it and its members', in agreeing with other people about it, and responsible for its usefulness in supporting successful inductive inferences. (p. 96).
- It would be question-begging to insist that, given your identification dispositions with respect to Swamp Man, and given the likely sameness of internal phenomenology between Swamp Man and you, he *must* have beliefs, intentions, etc. with the same kind of content as yours.

Ch. 7: Philosophical Analysis; Referents of Names

7.1 Overview

7.2 Philosophical Analysis (p.97-99)

- p. 98: "To know what a naming-term *token* refers to is merely to be able to same-track its referent, that is, to know how *sometimes* to recognize information about the same that is carried in other ways --- even just being able to recognize other tokens of the same name as such.
- Two q's: (1) (methodological) How do we investigate the nature of what a certain name names; (2) (logically prior) What constitutes a term having a certain thing as a referent. Re (1) – a negative point: we do not make progress simply by considering possible situations and deciding 'what we would say'.
- Negative claim re (1): Because names are directly referential and anchored to real kinds or categories, "imagining merely possible situations, even examining real but

bizarre ones, and trying to decide what one would say about them cannot help to reveal referents of actual names in actual language" (p.98).

7.3 Referents of Names (p.99-101)

- Q (2) is: "What does a name have to do to keep itself in circulation, to keep itself reproducing?"
- P. 99f.: "[A] name survives because speakers are using it and hearers are interpreting it in ways that produce agreement. Their uses and interpretations produce beliefs and expectations that are often enough confirmed and that lead infrequently enough to contradiction … On another and more fundamental level, however, the agreement hence the name's survival is accounted for in a Normal way by the fact that there is some one real individual or property or kind about which the name consistently carries information…".

AND it's "the individuality or the unity" of the underlying worldly 'anchor' – that which is being is being tracked -- "that has accounted for the fact that different people and the same person again, though perhaps using quite different methods to identify it, have found themselves in agreement about it" (p. 100).

- Dispositions to use the name in counterfactual situations are not relevant. And no particular information about a name's referent is passed along to language learners, nor is any particular method of identification attached to a name (see Ch. 12).

7.4 Theory Change in Science (p.101-102)

- A standard q in science: How could the extension of a term remain the same despite radical changes in ways of identifying it or in properties associated with its extension?
- There's no mystery if we take scientific terms to be directly referential terms for real kinds or properties. Same-tracking of these is the crucial function, and it can survive radical theory change.

7.5 Observation versus Theory (p.102-103)

- For Sellars and others, concepts were "nodes in an inference net". Nothing is 'given' directly to perception. So no sharp line between observation and theory.
- "If we deny that inferential processes are what *create* intentionality, substituting unicepts for concepts yields almost exactly the same results" (p. 102).

(1) "All observation judgments are fallible in principle."

(2) "Observation judgments can be made on the basis of signs of things without employing unicepts of those signs. So the doorbell might be heard directly as someone at the door, … the expression on Mother's face seen directly as Mother's angry state. Thus the distinction between what is observable and what is not is no longer absolute but relative to psychological processing" (p.103).

7.6 "Theory of Mind" (p.103-106)

- At issue: Does the use of language require speakers and hearers to represent one another's mental states, and does that in turn require a 'theory of mind'? Ch. 13 will

argue that the ability to represent mental states isn't needed for the emergence of language.

- Re the second q: although linguistic messages are sent purposefully and receivers are aware that they are, we should not overstate what this involves.
- Given the blurring of the theory/observation line (see above), we should realize that *representing someone's mental state* can be accomplished without having factic theoretical knowledge about the nature of mental states, and "recognition of another's mental state might involve merely affording knowledge rather than factic knowledge… One needs only to be able to recognize it [the mental state] in some way or ways so as either to collect information about it or to learn how to deal with it" (p. 104).
- We should reject the dichotomy of mind vs. behavior reading (as well as between unobservable/observable).
- Wasp example: the wasp doesn't have a unicept for the state of aggression, but it doesn't follow she only recognizes and represents black spots. The neural response mediating the wasp's response to the spots would be a pushmi-pullu representation. It "perceive(s) the current angry state of another wasp by observing its face just as a person may perceive the angry state of another person by observing their face" (p.106). But it perceives the angry state only as an affordance.
- Understanding another as 'goal directed' is recognizing that *something about* the other "lends it a versatile disposition to arrive at such and such an end state". This should be read simply as meaning that one has a representation or a unicept – probably one that recognizes this only as an affordance -- for what are in fact (real) purposive states, which "does not require knowing anything in particular about these states" (p. 106).
- "If a dog is designed or has learned to perceive symptoms of a squirrel's intention to escape up a tree and to react to such symptoms appropriately, … the dog's perception is a representation, presumably a PPR, of the squirrel's intention. It will produce its designed effects in accordance with a Normal explanation only if it indeed corresponds to such an intention. And if the dog is able over time to learn new ways to recognize when a squirrel is set to escape up a tree, then the dog has, at least an affording unicept and unitracker for squirrelish intentions to escape up trees." (p. 106)
- The key point: to have a representation of the relevant target (here: another's mental state) you don't need to grasp its "true nature. What makes it the case that the dog has a representation of the squirrel's purpose is that the dog's perception is 'locked onto' (through design by nature or by learning) to (what is in fact) the squirrel's purpose.

PART II: Infosigns, Intentional Signs, and Their Interpretation

Ch. 8: Introduction to Part II

(This chapter introduces Part II.)

8.1 Overview

8.2 Infosigns and Natural Information (p.109-111)

- The purpose of Part II is "to explain how unitrackers and cognition more generally use signs that carry … 'natural information'" (p. 109).

- 'Natural information' is carried by signs called 'infosigns.' All natural signs (smoke, rashes, energy patterns engaging the outer sense organs) and some intentional signs (some linguistic symbols, maps, charts and so forth) are infosigns. All and only those intentional signs that are true or correct are infosigns. Natural signs are neither right nor wrong; only intentional signs are true or false.
- All unitrackers have the same basic job – to collect information from infosigns.
- "An infosign token is… one member of an infosign-infosigned pair that exemplifies a non-accidental correlation between signs and states of affairs, the signs all corresponding to the infosignified states of affairs according to the same projection rules." (p. 110)
- The meaning of an infosign token is not determined by its sign-design alone but by its membership in a particular infosign family. A unitracker has to determine which *family* a sign design token belongs to, in order to discern whether it carries information that concerns the unitrackers target.
- p. 110: rash example.

8.3 Infosigns and Intentional Signs (111-112)

- Some, but not all, infosign tokens are also intentional signs. (A token rash is not an intentional sign.) And not all intentional signs tokens are infosigns. (A false intentional sign is not an infosign – at least not one that carries information that coincides with its intentional content.)
- Intentional signs are produced by mechanisms whose proper functions are to produce infosigns that "map onto the world according to rules rendering them readable by cooperating partners [of these mechanisms, that is,] by symbiotically designed interpreters or 'consumers'". These symbiotic producing mechanisms have emerged from a process of selection (genetic, perceptual tuning, learning, cultural). (p. 111)
- Like infosigns, intentional signs are neither determined to be such, nor determined in content, merely by their vehicle type. (They are individuated by their reproductively established families (para 12.3).

8.4 Interpreting Linguistic Signs (112-113)

- Ch. 9 & 10 will introduce many examples of infosigns. Next (Ch. 11) will come a description of intentional signs, which are "Normally produced by mechanisms whose function is to produce infosigns" (p. 113) This will be crucial for the (*naturalistic*) account of language to be developed. When all goes well – i.e. Normally – intentional signs are infosigns carrying natural information (Ch. 12).
- Key aspect of the account: replacing Grice's analysis of linguistic meaning with "one that makes no reference to speaker intentions". Linguistic interpretation is continuous with the interpretation of nonlinguistic infosigns and does not require recognition of speaker intentions. (Ch. 13)
- "Understanding language is translation of outer intentional signs into inner intentional signs" and is "a form of perceptual processing" (p. 113) (Ch. 14).
- Certain peculiar linguistic signs – ones that don't seem to submit to the above account – will be discussed in Chs. 15 & 16.

Ch. 9: Indexicals and Selfsigns

9.1 Overview

DB: A bit of background: In philosophy of language, indexicals and other 'context-dependent' linguistic expressions were for a long time seen as requiring special treatment, to be provided only after a 'purely semantic' theory for a language was laid out, which specified the fully conventional rules linking words and world. On a traditional view, indexicals' contribution to the truth-relevant meaning of a sentence could only be partially specified through conventional semantic rules (which rules could be thought as what one would master in virtue of learning the language). The rest was left to 'context', which meant either that it could not be theorized or that it required appealing to speakers' intentions/hearers' interpretations. The view laid out in this chapter is a radical departure from this tradition. First, RM proposes that we should ***begin*** *an account of language by thinking of indexicals. Secondly, she gives a* ***semantic*** *account, not relegating indexicals to pragmatics (and, along the way, deconstructing the sign/context separation). And, third, she understands their semantics as continuous with that of* ***informational signs*** *more generally, thus taking a first step toward treating all (Normally functioning) intentional signs as infosigns, whose (natural semantic) function is to carry information about bits of the world and whose interpretation does not require deciphering communicative intentions.*

9.2 Assumptions to be Questioned (p.114-115)

- p. 114: The chapter's project: "to go straight for the correspondence that exists between a true or fulfilled sentence token containing an indexical or demonstrative and the world affair that makes it true or fulfills it."
- Points of departure from familiar accounts:
- Part of the "context" as traditionally understood can sometimes be a "proper part of the conventional sign itself"; "not all meaningful elements of linguistic signs have *prior* meanings" (p.115)

- *Comment [DB]: Understood in terms of acquisition, this would seem to mean that an important part of acquiring a language is not a matter of mastering 'conventional meanings', but rather a matter of learning to use bits of the world to communicate.*

- Some signs are *selfsigns*: components that function as components of both the sign and the signified – whether by standing for themselves (in certain contexts) or by standing for things that bear certain internal relations to them.
- Distinguish 'absolute' and 'relative' reflexive signs.

9.3 Components of Conventional Linguistic Signs (115-117)

- We tend to think of the components of conventional linguistic forms as "fashioned by the signers themselves" – sounds, ink patterns, gestures. But such fashioning (even at the token level) is not intrinsic to the conventionality of language. (We can imagine a box of small blocks with symbols on them assembled variously to produce sentences.)
- ASL: illustrating how "things conveniently found lying around [including body parts] may be fully incorporated into the language itself, helping to complete fully conventional linguistic patterns". So "some of the things that were there" in the 'context' prior to linguistic production "may have been co-opted to help in composing

the conventional sign itself" (p. 116). And the co-opted elements can come from different sensory modalities. (Think of prosody.)

- So "the boundaries of a conventional mostly spoken language" need not "lie where the sound ends and the rest of the world begins". This theme will guide discussion of indexicals (p. 117).

9.4 Preliminary Examples of Selfsigning Components (117-120)

- "Spinach" example: a can has a label saying "Spinach" (p.117).

Claims:

(i)The word "spinach" by itself has no more meaning than the can does by itself.

(ii) Once the can is labeled, it becomes a significant part of a meaningful conventional sign, utilizing a certain conventional 'precedent for coordination'. (This is what opens up the possibility of a *mis*labeling: having a can of beans with the affixed label "spinach".)

(iii) The can of spinach stands for itself.

(iv) It is an illusion that in such cases no interpretation of the reflexive part (the component that stands for itself – the can of spinach) is needed. The distinction remains between sign and signified. It's just that a single thing serves as a component of both.

- Football team example (p.118)
- Stop sign example (p.118-119) – where the sign is placed is where one is to stop. So the place constitutes a component of the complete sign. With a red light: add time for *when* to stop.
- "Ouch!" (p.119): the person who cries out stands for herself, with "ouch" serving as a label on her indicating that it's *that person* who hurts. Shaking my fist is producing a conventional sign that I, the one shaking the fist, is angry.
- "It's raining" (p.119): by default, the time and place are those of the utterance. (Not so for past tense sentence tokens. What needs to be fixed is the relevant relation to the time of the sign (concomitant, before, later).
- In English, directives given in the imperative mood, time and addressee of speaking are both used as part of the sign.

9.5 Indexicals and Demonstratives (p.120-122)

- Assuming that "whatever is necessary to observe in relation to a conventional sign in order to grasp its truth conditions, granted that one knows all the relevant conventions, must be part of the sign" *(see previous comment)* the addressee of "Please go!", like that of the "you" in "Would you please go!" will be an actual part of the complete sign.
- But now "You" in "Would you please go!" can be seen as simply an *anaphor*. "Its contribution is to make explicit the grammatical place of the addressee as a sign component within the grammar of the sentence, hence the place of the addressee as a signified component within the state of affairs represented" (p. 120).
- The so-called indexical is a "dummy anaphoric 'pro-word'" that is inserted into the relevant syntactic position to show the position of the demonstrated object within the represented state of affairs.

- The hearer needs to figure out how to *anchor* the anaphor. This can be achieved by various conventional – and sometimes nonconventional – means. (Conventional: e.g., gender and number marking, adjectives). But, whereas on the traditional view, 'pure demonstratives' require pragmatic interpretation, on the present view, demonstrative pro-words are sometimes anchored to their 'intermodal antecedents' conventionally. (p. 121).

9.6 Addendum on Intensional Contexts (p.122-123)

- A take on Davidson's analysis of intensional contexts: Davidson: "X said that S" □ "X said *that*: S." "Galileo said that the earth moves." □ "Galileo said that: The earth moves."
- Two separate sentences, the first demonstrating the second. Critics: S is not a detachable part of the original sentence.
- Gloss: Both Davidson and the critics are right. S is a proper part of the original sentence, *but* it's a part that refers to itself "doubling as a sign and a signed component". This has the advantage of fitting in with a general perspective on language in which selfsigning is the rule rather than the exception.
- Qualification: In intensional contexts it's only the *representational content* of the embedded sentence that stands for itself (allowing for variation in reporting and reported speech/thought). Then it is the content that is, as it were, 'held up for display', standing for itself. Alternatively, it may be the words used that stand for themselves as in "Peter thinks that woodchucks are bigger than groundhogs". That is, Peter thinks that the things called "woodchucks" are bigger than the things called "groundhogs."

Ch. 10: An Anatomy of Signs

10.1 Overview

10.2 The Project (p.124-125)

- 2 central q's:
  (i) By what rule of correspondence does a sign map onto its signified?
  (ii) What is it for one state of affairs to *be* an infosign for another? (Ch. 11)
  - *Gloss (DB): (i) is a semantic q, concerning the right assignment of 'meaning' to a sign. (ii) is a 'metasemantic' or constitutive q concerning the signing relation – the conditions that make it the case that a sign is an infosign.*)

10.3 Infosigns Are Always Articulate and Often Productive (p.125-127)

- An infosign is not: the sign design a of sign, a token of a sign, and so on, taken by themselves. Nor is it a sign design. It is a particular, situated in our space-time causal order.
- "[A]s represented, states of affairs are always articulated into components or aspects." (p. 125)
- The representations – the infosigns – are all like complete sentences in that they are articulated into elements/aspects that represent corresponding elements in the

state of affairs. The elements/aspects can be removed and different ones substituted in one by one to yield representations of alternative states of affairs.
- (Notice that that is more restrictive than being just "recombined.") Infosigns fall into *systems*. This makes for a certain kind of systematicity and productivity (contra common prejudice concerning the uniqueness of human language in this regard).
- Infosigns are always articulated into a variable and an invariant component, where variable components can be replaced via what mathematicians call 'substitutition transformations' (not just "recombined," producing signs that correspond systematically to corresponding substitutions in their signified states of affairs. This applies to natural signs, as well. (The sound entering your ear of a dog bark can be an infosign with variable concerning size, age, direction and distance, and current mood.)

10.4 Reading Infosigns (p.127-128)

- Given the (semantic, non-formal) individuation of infosigns, what meaning a sign carries is not carried on its sleeve – it has to be determined by context and 'pedigree'. (An epistemological problem.)

10.5 Infosign Systems and Families (p.128-129)

- "Infosigns always come in systems. … Tokens from the same infosign system belong to the same 'infosign family'." (p. 128) – e.g. "black clouds that mean rain (…), animal tracks (…). Correctly executed bee dances (…). Mercator maps having the same keys, … sentences in the same language that have the same syntactic structure…". "A cognitive system that is tuned to read one sign from an infosign family is tuned to read others."
- A glaring exception: infosign families that contain variables filled by *names*. (See later, 10.7.4.)

10.6 Variants and Invariants; Embedded Infosign Families (p.129-131)

- Infosigns in a family will be articulated into an invariant and one or more variant components/aspects, where different values for the variables will yield different family members. (Bee dance example; blueberries example.)
- Re: the interpretation of infosigns: "The strength of an infosign in a context is a matter of the strength of its family within that context, a matter of statistics on the frequency with which surface-format tokens matching the family are indeed family members" (p. 130).
- Infosign families may cross over other families – depending on what are taken to be their invariants.

10.7 A Taxonomy of Infosign Variables (p.131-136)

- There is a surprising diversity of correspondence relations between variable infosign components and their signifieds. These are "correspondence rules that take one from a sign element to what it signifies" (131).

10.7.1 Absolute reflexives: 'selfsigns' (p.131)

- "The simplest infosign variables take reflexive signs, selfsigns, as values." – e.g. "a goose performing a mating dance to sign that he himself is ready to mate". (p. 131)

10.7.2 Relative reflexives (p.131-132)

- Relative reflexives elements have their signifieds determined as bearing a certain relation to themselves (e.g. a certain number of days later than this day, a certain multiple of this length).

10.7.3 Isosigns (p.132-133)

- Arguments for *isosign* variables come from *different* natural domains than their signifieds but from domains that are isomorphic. Their signifieds are determined by a certain mapping from one domain to the other.
- (Graph ex.). "[T]he height of a certain bar on a certain graph may be an isosign standing for the size of the profit of a certain company in a certain month" (132).
- Isosigns can be analog or digital components of infosigns. The more dimensions (especially analog) mapped by an infosign, the greater the temptation to think of the sign as a picture-like rather than sentence-like (think of perceptual representations). But even perceptual representations involve elements that are names, rather than just isosign or reflexive sign elements.

10.7.4 Names (p.133-134)

- "Names are values whose signifieds are not derivable by any general rule from their sign designs." The meanings of names are by and large 'arbitrary' and have to be learned individually.
- Examples of names/naming components: different colors, or different line types, showing different types of roads or showing different data sets. Symbols for cities, mountains etc. on a map. Simple numerals. Words that designate individuals, kinds, properties, and so on, in natural language. Specific places, times, and shapes (think of these as being like labels on things telling what is in them).

10.7.5 Roaming names (p.134-135)

- Most names in languages are "roaming names" – they move freely across constructions and places in constructions, filling variables in signs from many different infosign families. ("John has a cat", "cats" in "Cats are smarter than dogs" (different syntax) or "please" in"John is eager to please", "John is easy to please".) Roaming is central to the individuation of languages and greatly facilitates human unitracker development. It is made possible by properties of the typical linguistic media (speech, print, gesture).

10.7.6 A summary example (p.135-136)

- The gas gauge example. "It is not only language that is articulate and productive" (136).

10.8 How Language Is Put Together (p.136)

- The starting point for understanding how language works should not be contextless roaming names that could serve as [*atomic semantic building block*s]. "It is crucial … to see that all representation, outer or inner, is, at root, representation of full states of affairs by full states of affairs. Words signify only when they are embedded in sentences. Like dog barks, cans of spinach, and times, they have meanings only in that they can come to signify in certain contexts" (p. 136).

Ch. 11: Infosigns and Natural Information

11.1 Overview

11.2 The Project (p. 137-138)

- The wider project is "to find a description of cognition and a sense of information such that we can understand how cognizing animals can learn from signs that bear information. We want to understand how there is a kind of sign that carries a kind of information that makes cognition … possible" (p. 138).
- The present chapter aims to explain what is distinctive in common to so-called 'natural signs' – nonintentional, infosigns.

11.3 Examples of Nonintentional Infosigns (p.139-140)

- Given the great variety of infosigns, it is a challenge to identify central features that "can explain how it is that these signs can support induction and cognition".
- p. 139: A list of 12 such examples. [1. black clouds, 2. fever, 3. boiling water, 4. voice quality, 5. shape of head, 6. first part of a song 7. going to the party, 8. diamond shape, 9. Polaris, 10. magnetosome, 11. passing a pond, 12. Suzy's mitten.]

11.4 Causal Connections: Dretske on Natural Information (p.140-141)

- Although mediating causal connections of some sort between sign and signed (in either directions) are involved in many cases of usable signs, they are neither necessary not sufficient for the production of the relevant relation (see examples).
- Dretske's 'natural information': does not admit a privative form (no misinformation) ; identifies signs with 'sign designs' that always correspond, by natural necessity, to a given type of signified.
- As Dretske recognized, there is never an objective probability of 1 that a sign (understood as a surface format) will correspond to a certain signified without 'channel conditions' being in place.
- In general, animal cognition is not underwritten by infallible indicators. Signs used in actual perception/cognition don't carry their infosign types on their sleeves.

11.5 Correlational Information (p.141-143)

- On a more realistic picture, perception/cognition relies on mere natural correlations rather than necessities. "Correlational information" is produced "when there is a *non-accidentally continuing* correlation between one kind of state of affairs [A] and another [B], such that the probability of the one kind is raised given the other" (p. 141).
- The non-accidental correlation can be of various types – not necessarily causal. We only need to insist that (i) an A-affair token doesn't carry the information unless a corresponding B-affair token actually exists ('no misinformation'); and (ii) the correspondence is due to the same reason that the correlation continues more generally ('non-accidental').
- A strength of this characterization is that it admits of any kind of non-accidentally continuing correlation between two kinds of states of affairs. (Note that sometimes we have "indirect correlational signs" – they too can count as signs. E.g. Suzy's mitten.)

11.6 The Reference Class Problem (p.143-144)

- Difficulties for the correlational description of natural signs:
  (a) The 'strength of correlation problem' (to be discussed later):
  (b) The 'reference class problem': Which of the innumerable possible reference classes is the relevant one for grounding a correlational sign?

11.7 Addressing the Reference Class Problem (p.144-147)

- Drestke's notion of information: an objective commodity that requires/ presupposes *no* interpretive processes, vs. Peirce's idea that built into every sign an 'interpretant'. A compromise: "information as something that can be found out in the world prior to any actual interpretive processes but that is defined by its capacity to be interpreted" (p. 144).
- Signs as *food for cognition*.
- To serve this purpose, the correlations that underwrite signs must be found in the space-time paths of the animals that discover and use them.
- "Because that something is a sign reasonably depends on the possibility of its being used as a sign (…), and because whether it can be used as a sign is relative to the range and the capacities of the animal that would use it, whether it is a sign … is best understood as being relative to an animal, species, or society rather than being an absolute matter" (p. 146)
    - (*So this is a sense in which, although information is objectively found in the world, it is interpreter-dependent; hence the compromise between Dretske and Peirce.*)

- Just as what can serve as food depends on the peculiarities of the animal's digestive system, so what can serve as a sign depends on the peculiarities of the animal's cognitive system.

11.8 Using Infosigns (p.147-148)

- Serving as a sign for an animal requires that the animal (1) lives within the context of the relevant correlation, (2) has systems capable of recognizing the sign designs, (3) can collect enough examples of the correlation, (4) can read other signs that might mark the boundaries of relevant classes.
- In the case of humans, given the role of cultural learning of data collection, there is no way to determine which kinds of correlations might serve as infosigns. So 'natural information' remains an open and pragmatic notion.

11.9 Correlations between Types of States of Affairs (p.149-150)

- Strictly speaking, correlations obtain between, e.g., black clouds appearing at a time and place and rain falling shortly thereafter at that place (not simply "between black clouds and rain"). And the correlation has to be such that the one state of affairs is in a certain historical relation to the other.
- It follows that "the correspondence rules for *intentional* signs (which include linguistic signs) necessarily correspond to correlations between past actual signs and past *actual* signified states of affairs of the relevant kinds … [so] for every such family of intentional signs there must exist a corresponding subfamily of infosigns" (p. 149)

- An infosign must correspond to its signified *for the reason that generated the correlation*. Accidental pairings are not infosign-signified relations.

11.10 Infosign Strength and Response Strength (p.150-151)

- Infosigns come in different strengths. "The strength of an infosign token for an animal is the strength of the correlation within the reference class that the animal relies on when trying to interpret the token. It is a measure of how easy it is for the animal to tell that a sign of X is indeed a sign of X" (p. 150)
- Accordingly, an animal's interpretive response to a given sign design may also come in different strengths.
- Returning to the issue of strength of correlation (11.5), given that correlations of *any* strength can support a serviceable sign if the consequences of failure to respond to the sign are disastrous enough, there's no minimum strength that every correlation must have in order to ground an infosign.

11.11 Redundancy (p.151-152)

- Sometimes more tokens of the very same infosign increases the strength of the tokens taken together (e.g. many deer's hoof-shaped prints near one another in the mud). Many different kinds of infosign which can be read as signs of the same thing may greatly increase the strength of each. The reliability of the human perceptual systems is greatly enhanced by exploiting redundancy.
- As mentioned in 5.9, redundancy also serves as evidence for the general adequacy of one's unitrackers. "That one often seems to be able to discover the very same things to be true in a wide variety of different ways is good evidence that one is tracking real entities in one's distal world" (p.152).

11.12 Metacorrelations (p.152-153)

- A single instance of a correlation can generate knowledge of it given suitable knowledge of metacorrelations over real categories of real kinds.
- "[B]y knowing what kinds of features tend to generalize over kinds in the same category, one can often learn a great deal about a real kind by examining just one member with care" (p. 153).
- Some metacorrelations predict regularities based on the enduring of individual objects.

11.13 Indirect Infosigns; Indirect Natural Information (p.153-154)

- Suzy's mitten functions as an *indirect infosign*; it exploits an ability to gather indirect information.
- The case of Bluster, the bomb maker: he has calculated that the fizzing would be an infosign of imminent explosion, given a chain of (causal) correlations.
- Whether or not an infosign is indirect depends on the statistics. (The fizzing might, of course, be used as a direct sign of impending explosion, supposing that a sufficiently large number of bombs of this kind have been made and their action observed.)

Ch. 12: Intentional Signs

12.1 Overview

12.2 The Project (p.155-156)

- The project here: "to understand how coming to knowledge of the world through language is but one more way of coming to knowledge of the world through infosigns, signs of essentially the same kind that unaided nature often provides within the natural world" (p. 155).
- Recall: intentional signs are signs that can be *false* (or unsatisfied); so we need to understand what about this class of signs makes them liable to being false. (Intentional signs form a real kind and there's no expectation that we can provide necessary and sufficient conditions.)
- We begin by describing an intentional infosign family that is composed of just those tokens of a certain infosign where intentional content coincides with informational content. These are intentional signs that are true.

12.3 Intentional Infosigns (p.156-159)

- p. 156: An intentional infosign is one member of a sign-signified pair which pair comes from a family of like pairs that has continued to exist because:
(1) instances of the sign-signified correlation/pair form a reproductively established family (REF), roughly, a family whose members are alike because copied or reproduced from one another.
(2) many members of this REF perform a communicative function serving a sender and a receiver (themselves designed to communicate using the pair)
(3) the pairing continues to be reproduced for that reason
(4) resulting in a non-accidental correlation that constitutes the infosign family.
- The crucial aspect: *however* things got started, the continuation of the correlation is an instance of symbiotic adaptation: it depends on its usefulness and its usefulness depends on the correlation. (Newer tokens of the infosign type map onto their signifieds *because* older tokens did, thereby producing an effect that causes reproduction of the mapping.)
- Senders and receivers using intentional infosigns have been tailored by learning or natural selection to cooperate with one another through the use of these signs. As with other symbiotic relations, it's the traits of the partners, here the senders and receivers, that are selected for or stabilized so as to fit one another.
- Effecting communication between sender and receiver pair is (something very like) the proper function of intentional infosigns.
- The shared purpose of communication via intentional infosigns can be common to sender and receiver because useful to each separately. (The cat stretching toward the door to go out.) Or due to kin selection. (Rabbits' danger thump.) Or because what benefits one automatically benefits the other. (As with inner representations within an organism.) Or because of the interests of one partner to aid in the project of the other. (Military example.)
- In the case of humans, it appears that there is an innate tendency to discern and adopt others' purposes and dispositions to help them, which yields an impulse to communicate information cooperatively. (Cf. Tomasello.)
- Fn 3: Dispositions to cooperate are not, as such, dispositions to altruism.

- [p. 158f.: On the distinction between direct and indirect reproduction; membership in higher order REF.]

12.4 Intentional Signs and Stabilizing Functions (p.159-160)

- "A sign design token is an intentional sign (though maybe not also an infosign) if it is a member of a REF of sign designs that originated with the designs of an intentional infosign-signified REF" (p. 159). (A merely nervous rabbit thump, "It's raining" said when it isn't, are intentional signs, but not intentional infosigns.)
- Intentional signs as such belong to REFs that have proper functions. Descriptives have the basic stabilizing (or symbiotic) function to prompt or remind of true beliefs in hearers; directives have the function of prompting indicated hearer actions. (The functions are only performed Normally when the signs are infosigns.)
- The 'intentional content' of an intentional sign is determined by the mapping function that is distinctive of the REF from which it is derived.

12.5 Pure and Impure Intentional Signs (p.160-161)

- A doe's abrupt freezing when she senses danger may be a nonintentional infosign that is useful to fawns in alerting them of danger. If the deer's way of freezing and the sensitivity of her fawn have been selected for partly because they serve this function, then it is also partly an intentional sign, an *impure* intentional sign.
- By contrast, a *purely* intentional sign is one where reproduction of infosign-signified pair depends entirely on repetition of its communicative function – it has no function beyond communication.
- (p. 161: Sitting down at a restaurant as an example of an impure intentional sign, a request that the waiter serve you. Extending one's hand toward another.)

12.6 Entwining of Intentional Content and Nonintentional Information (p.161-163)

- The doe's abrupt movement can become exaggerated, to make it easier to read. It makes no difference to the fawn whether the sign is intentional or not. Intentionality concerns only the reason for the correlation grounding the infosign. Involuntary screams can be intentional calls for help designed to bring others to one's aid or merely to scare attackers. All that matters from the interpreter's perspective is that the scream is an infosign of needing help.
- A Normally true intentional infosign will often carry considerable nonintentional information – e.g. the location of the calling vervet, or the beliefs of a speaker.

12.7 Intentional Signs Used by Non-Human Animals (p.163)

- Communicative signals used by non-human animals are intentional signs, mostly designed through natural selection, but some developed through learning.
- They all seem to be PPRs, so location reflexive, having time and place as self-signing elements, but at times also the signaling animal itself. (The mating dance of Greygoose is a sign that Greygoose himself is ready to mate.)

12.8 Maps, Charts, Diagrams, Graphs (p.163-164)

- Each sign making up a map is an infosign of the features or domain mapped. More 'local' mapping relations (between bits of the map and places) will not have been reproduced owing to *functions* they have served. "It is the whole of key-plus-

adherence-to-key that constitutes a member of the relevant … REF". So it's only the map taken together with its key that can misrepresent, strictly speaking.

12.9 Extending the Senses (p.165-165)

- Various human-made instruments and machines are designed to produce infosigns to be interpretable by trained and equipped observers; they have the production of these infosigns as their proper functions, derived from makers' intentions. The produced infosigns are intentional signs.

12.10 Inner Representations (p.165-166)

- It is plausible to suppose that well-functioning cognitive systems of individual animals use inner intentional infosigns of its environment to help guide its behavior, with initial producers being the perceptual systems and consumers being faculties of higher cognition (in some cases) and ultimately the motor systems, with the most basic inner representations being PPRs.
- As a side effect, there may be a lot of natural information produced that the animal's cognitive system does not know how to recognize and use (hence is not intentional content). Yet the consumer systems of humans may learn to interpret this additional information. (And this may be important for understanding differences between us and other animals.)

Ch. 13: Linguistic Signs

13.1 Overview

13.2 The Topic (p.167-170)

- When studying language, as when studying human physiology, understanding Normal function – what accounts for survival and spread – is logically prior to understanding deviations and malfunction.
- P. 168: When language functions Normally, a transfer of mental states occur from person to person (descriptives transfer beliefs; directives transfer intentions/purposes; identities transfer dispositions to coidentify; conditionals – dispositions to acquire attitudes expressed by the consequents upon having/acquiring those expressed by the antecedents; normative forms transfer attitudes and modals transfer modal attitudes).
- But it is more accurate to say that the function of e.g. a descriptive is to transfer a *true* belief (compare: the function of the esophagus is to transfer food, not just whatever is in one's mouth, to one's stomach).
- Priority of the descriptive and directive: other functions are parasitic on their functions (e.g. the capacity for believing something possible is parasitic on the capacity for thinking something true).
- PPRs again: we can think of intentions, like representations of affordances, as two-faced PPRs. They are representations, in this case, that tell at once *what we are to do* (a directive face) hence what the case *will* be (a future-descriptive fact).
- They differ from perceptual representations in that they map variations in goals (rather than the perceived world) onto represented future states of the world (rather than onto the organism's (possible) actions. But in the case of intentions, the contents of the two aspects coincide).

[e.g. If I intend to open the door, I harbor a representation that has the directive aspect *that the door is open*, which content also captures the content of the future-descriptive aspect – how the world will be.]

- Directives → descriptives: if forming an explicit intention involves forming a belief about your future, then understanding directives involves the same translation process (from outer to inner signs) as would understanding future descriptives. So the study of directives can be subsumed under the study of how descriptives are constructed and understood.

13.3 Semantic Meaning (p.170-173)

- What makes a rule 'conventional'? What makes some uniformity a *rule*? Proposal (1988a): Linguistic conventions are 'simple conventions' – "patterns of speaker-hearer behavior that are handed down or reproduced" – designed to solve coordination problems (in Lewis's sense) – they "coordinate speaker utterances with hearer responses so as to serve purposes of both" (p.170).
- Semantics describes the conventional side of language – the specific proper functions of various linguistic forms (that which explains continued reproduction by speakers and continued uniform responses by hearers).
- Two parts: the *proper* or *stabilizing* function of linguistic forms (their role in communicating attitudes, motivating others and so forth, and the *correspondence rules* (which constitute 'narrow' semantic meaning – that which derives directly from intentional content).
- p. 171: The correspondence 'rules of language' are "uniformities in use that result from speakers and hearers mostly agreeing, given what actually turns up to talk about in their environments, on what it is to 'go on as we did before'" (Lewis, §2.3).
- Re the semantics/pragmatics distinction: On Korta and Perry's (2015) classification, meanings divide into

(1) *semantic meaning* as carried by tokens of a type regardless of context (addressed by Gricean semantics)
(2) *'what is said',* which is meaning with contextual elements fixed
[This is what Morris and Carnap considered to be the domain of pragmatics]
(3) *implicated meaning* as fixed by speakers' hearer-directed intentions
(addressed by Gricean pragmatics)

- [DB: For Grice, (1) – 'sentence meaning' – is constructed out of regularities in *speaker meaning*, which is in turn determined by ('occasion') hearer-directed speaker intentions. Such intentions can also generate implicated meaning – as in (3) – even once sentence meaning gets stabilized by convention.]

The view defended here treats (2) as an aspect of (1).

- *[This is in keeping with what has been described as 'semantic pragmatics'.]*

- The reference of names, indexicals, demonstratives is fixed *not* by intentions but by the intentional sign REFs (reproductively established families) from which they were copied. You can't *make* a token of "Bill" refer to John by intending it to (or by 'having John in mind' when producing it); and if your token of "Bill" was copied from previous tokens used to refer to Bill Brown, then you said something about him, and not about Bill Green. Thus, 'what is said' is fully given by the semantics of the language.
- Semantics thus understood individuates expression types not in terms of surface form (or 'sign design') but rather in terms of the constructions they represent.

- Constructions are the REF's formed by copies and copies of copies of a certain meaningful expression. (e.g. "Hit me!" used in Blackjack to request a new card is a token of a different construction from "Hit me!" used to request an assault. Constructions – single words, phrases, sentence frames – are typed in terms of the REFs to which they belong. But often they can only be told apart using context.

13.4 Communicating with Language: The Broader Picture (p.173-175)

- "The function of descriptive language is to supply natural information to hearers/interpreters. Language would not survive if speakers did not speak the truth enough of the time. True descriptive language carries natural information, exactly as nonintentional infosigns carry information. The simple hypothesis, then, is that this language is interpreted in the same way that nonintentional infosigns are interpreted. Interpreting descriptive language in use is just one part of the wider project of trying to discover what's in the distal world by interpreting, ultimately, infosigns ruffling the sensory surfaces" (p.173)
- This is consistent with the fact that *sometimes* hearers acquire information from language by treating it as a sign for the speakers' beliefs (and taking into account other meta-facts). But this isn't a necessary part of Normal language function. It isn't part of the basic functions on which the survival of language depends.
- "In fully Normal cases, descriptive sentence tokens are true, not accidentally but in accordance with a Normal explanation, and correctly understood. The speaker purposes to make an infosign of something for the hearer with the aid of a language known to both … The hearer purposes to read this infosign. If both speaker and hearer succeed, the hearer reads the infosign as a sign that at least includes what the speaker purposes it as a sign of … How does the speaker approach his task? The speaker constructs a sign that, in the context, means to him what it is his purpose convey. …The language functions Normally when hearer content includes the part of speaker content that the speaker purposefully put into his sign." (p.174)
    - [*Question (DB): How can this be accomplished without the speaker having thoughts about the hearer's (likely) thoughts?*
      RGM: Important here is to read *purposing* not as involving a mental representation of what is to occur but more broadly. Compare the purpose of your eyeblink reflex, of swallowing as you drink your coffee, of putting you left leg forward as you walk. Nor are "unconscious purposes" representations of goals that are furtively covered up.]

13.4 Meta-Regularities and Extra-Semantic Infocontent (p.175-176)

- To be understood, speakers must compose perspicuous linguistic infosigns, doing 'what has been done before'. But what guarantees that the hearer will understand this in the same way as the speaker?
- Uniformity in what are taken to be precedents might be explained if there were universal 'filters' in grammar and lexicon that constrained human dispositions to 'go on', linguistically speaking. A simpler explanation is that learners come to discern *meta-regularities* governing the lower level regularities producing correlations that ground infosigns. (E.g. discerning the pattern that names stand for natural units, individuals, real kinds, rather than for stages of individuals and mere classes. Once that pattern is discerned [presumably very early] lexical learning becomes very rapid).

- Connecting back to lexical meaning: to use "horse" as others do, there's no need to share any particular information about horses. The information we each associated with horses (some it having been acquired through hearing "horse") is *'extra-semantic infocontent'*, which is not, strictly speaking *communicated* from speaker to hearer but rather filled in individually. Purely semantic content is highly *skeletal*.
  - [*Question (DB) re last sentence of subsection – "Merely informational content rather than semantic content must be the main stuff language carries to hearers." ?? Is "carries to hearers" to be contrasted here with "communicated from speakers to hearers"?*
    RGM: yes.]
  - [*DB: A potential predecessor of this view in connection with the meaning of proper names is perhaps Searle's view, on which proper names are 'pegs' on which language users hang information they collect concerning the name bearer, where no particular piece of information constitutes part of the semantic meaning of the name, nor is it necessary to share information to communicate using the name, or to transmit any particular information in passing it on to learners. Neither semantic competence with a name nor successful communication with it require sharing of specific information about its bearer. RM extends this to all names, and extends the category of names to include not only individuals, but substances, events, properties, etc.*
    RGM: Yes. = "direct reference" as in §2.4]

13.6 Grice's Conversational Maxims (p.176-178)

- Much of the information imparted through linguistic utterances will be idiosyncratic. "Speaker and hearer will each bring their own knowledge of the world and their own interests to interpretation of the sign. … How, then, is successful communication achieved?"
- Part of the answer is that speakers and hearers typically share background information. Another part is to do with what Grice termed 'conversational maxims', the spirit behind which "is the very foundation of human language" (p.177).
- There is a "peculiarly human impulse to cooperate and hence to share information as needed". The underlying mechanism at work in linguistic communication – when effected in the Normal way – begins with joint attention.
- "*Joint* attention is not, of course, attending to one another's mind. It *is attending to the same thing as the other*… by following cues from the other's speech, from where they are looking, by understanding at what their current activity aims, and so forth. … *Shared attention* bends the speaker and the hearer toward noticing the same things and interpreting signs as of the same things." (p. 177, emph. added)
[RGM: DB says I am not using the term "joint attention" here in the common current way since I don't take it to involve theory of mind.]
- Explicit awareness of each other's mind, background, etc. is *not* part of the *basic* mechanism of communication. Not should it be thought to have been necessary for the evolution of human language.

13.7 Far-Side Pragmatic Meaning, or Semantic Meaning? (p.178-181)

- Some phrases have a single semantic meaning, others serve as sign designs for more than one construction, or contain tokens from several constructions. (One

cannot tell from the surface form whether a particular token of "dyed in the wool" exhibits three constructions or one).
- To clarify the semantic question regarding a particular surface form, we should ask whether (1) its tokens all belong to the same REF (through copying or superimposing) or (2) different tokens have different origins (e.g. "broke a finger" – one's own or someone else's).
- Contra Grice, usage-based theories maintain that there is no good reason to suppose that language learning is better explained by postulating that what is learned is a small lexicon plus syntax (how parts of this lexicon may be combined) used in accordance with pragmatic principles to yield contextual interpretations. (The opposition here is between generative grammars, which postulate autonomous syntax and compositional semantics and usage-based theories, which emphasize the ubiquity of 'formulaic' constructions, 'holophrases', 'chunking', and idioms; see fn. 8.)
- The claim is that "much more can be said in purely semantic, conventional, although idiomatic ways, hence much less may need to be improvised through pragmatics than classical Gricean pragmatics supposed." (p. 181)
- A third possibility [deconstructing the semantics/pragmatics distinction]: "Many constructions may be REFs of impure intentional signs. Their surface forms are sometimes reproduced whole and other times compositionally, these reproductions being linked, however, into a single REF network. They may be understood compositionally by some speaker-hearers and more holistically by others, half in half by still others. This would probably render the semantics/pragmatics distinction indefinite, in various degrees, over a significant portion of any language." (p. 181)

13.8 Addendum: Gricean Temptations (181-183)

- In Normal linguistic communication, there need not be a transfer of beliefs from mind to mind, but rather "a speaker's displaying to a hearer some piece of the world, … in the way things may be displayed for ordinary perception".
- Contrary to the Gricean picture, "[t]he hearer is not usually interested in what is in the speaker's mind any more than he is in what occurs inside the TV or the camera. He wants information about the world, not the speaker."
- Awareness of speaker intentions in speaking is a different kind of ability than the ability needed for Normal language interpretation (which is the ability to read – directly – natural information content carried by linguistic signs).
- Grice argued fallaciously from the premise that a hearer who *believed that a speaker did not intend* to produce in her a certain belief/intention would not acquire that belief/intention to the conclusion that a hearer who *failed to believe that a speaker intended* to produce in her a certain belief/intention would not acquire that belief/intention.
- On a 'transparent reading', "understanding what the speaker intends/means", the hearer must understand the utterance to have the same meaning the speaker understands it to have – which is indeed required for Normal language function.
- On an 'opaque' reading, it means that hearer must understand *that the speaker has a certain intention* and recognize what that intention is (and so be able to think about the speaker's intentions under that description). This is *not* required for Normal language function and is not a usual part of *normal* communication.

Ch. 14: Perception, Especially Perception through Language

14.1 Overview

14.2 The Project (p.184-185)

- The Normal interpretation of linguistic signs, like that of all infosigns, "begins with the reception of sensory data" carrying natural informational content "and ends with an understanding of that content". A stronger claim: "understanding language, when you believe what you hear as you hear it, is a form of 'direct perception' of what language is about, or that it is a form of 'perceptual understanding', meaning by an 'understanding' a representation employing unicepts" (p.185).
- Learning words is acquiring new unicepts and unitrackers.

14.3 Perception as Sign Reading (p.185-186)

- "Perception … attempts *translation* of the natural informational content carried by patterns in sensory data into inner intentional signs, either beliefs or representations of affordances." (p. 185) This is in contrast with *inferential* views of perception, which take *visual* perception, specifically, to involve preconscious abductive inferences from light patterns on the retina to their immediate causes in the light reflecting surfaces.
- An assumption here has to be that the perceptual systems make use of the laws of optics that are involved. Otherwise they could not obtain the premises required for such abductive inferences.
- Mere translation doesn't require inferences from effects to causes. What must be taken account of for interpretation is only the *semantic* relation between the sign and the signified, not the mechanism that accounts for it.
- The objects identified through perceptual processing need not be direct causes of the sensory data that serve as input for it.

14.4 Attached and Detached Signs, with an aside on Animal Cognition (p.187-190)

- Whereas perception is paradigmatically of the immediate environment, linguistic understanding can pertain to anything anywhere.
- 'Attached signs' are signs that show the space-time relation of their signifieds to the perceiver, e.g. patterns on the sensory surfaces that show not just objects but where these objects are relative to the perceiver, e,g.,black clouds or doorbells seen/heard from here now, road curve signs. 'Detached signs' are signs that do not show the space-time relation of their signifieds to the perceiver, e.g.,almost all linguistic signs.
- The capacity to coidentify elements of various different detached signs so as to understand what is/did/will happen in portions of the world detached from our present one is the hallmark of human understanding.
- Purely factual information can only be *stored* in inner representations that are detached. To store info about a past situation (for future use in action) the representation must be factive and use portable representations of place and time (e.g.,unicepts) rather than place and time themselves to represent place and time. "Knowledge of facts cannot, alone, guide action. Action requires grasping affordances, that is, understanding what to do or what can be done from here-now. Purely factual knowledge is useful only when combined with perceive affordances, with procedural knowledge." (p. 188)

- Re human vs. nonhuman cognition: “so much that is comprehended by humans but not by non-human animals seems to depend on reading detached signs that it seems possible this may mark a basic divide between them. Perhaps other animals don’t employ factic inner representations, thus cannot interpret detached representations, and for this reason can’t understand language as we know it. Perhaps all of their knowledge is procedural. Their behavior is guided by perception of current affordances, and by anticipation of affordances that will result if these current affordances are followed.” (p. 188f.)
- Communication among and with animals seems to be limited to attached signs (vervet monkey calls, beaver danger splashes, mating dances, Alex the parrot’s responses to “What color?”).
- Despite the differences between attached and detached signs, no reason to suppose the psychological process required to interpret them is different. No reason to suppose that ‘(uniceptual) inference’ is involved in one but not the other.

14.5 Genuine Perception Thought of as Certain (p.190-193)

- The facticity of verbs of perception (“I see/hear that p” entails “p”) does not mean that perception is always veridical, though it does reflect its relative reliability.
- The possibility of nonveridical perceptions (including persistent and transient perceptual illusions) may encourage a radically inferential view, according to which all we directly perceive are sense data; the rest must be derived by inference.
[So facticity doesn’t entail infallibility and fallibility doesn’t entail inferential processing. I can see the barn by seeing one of its walls, even though it’s possible for me to mistake a barn façade for the barn. If it’s only a façade then of course I didn’t see the barn, but that doesn’t mean that, if I did see the barn, I must have successfully *inferred* the presence of the barn from the wall I perceived. Along these lines, we can reject a standard argument that we can never directly perceive emotions, but can only infer their presence from behavioral symptoms.]
- The alternative is to accept the reality of perceptual illusion: “the methods used in translating infosigns that are present in sensory data directly into cognitive understanding are fallible.” (p.192) One can hear someone at the door (if there is someone knocking at the door) - but one can also experience a genuine perceptual illusion, incorrectly thinking one has heard someone at the door. A peculiar kind of sound design may be an indirect infosign of someone at the door.
- Perception understood as translation from infosigns must be fallible from the start, and not just when unicepts are applied. “[A] translation from infosigns must be superbly sensitive to infosign strength” (p.193).

14.6 The Contents of Perceptual Experience (p.193-196)

- A common assumption: all knowledge comes from *experience*, so any change in what is perceived must be implemented by a change in perceptual experience. Another assumption: each sense has a proprietary subject matter.
- From these it is inferred that “strictly speaking, no one is able to see that Peter is unhappy, smell that bacon is frying, hear that the wind is blowing, … even hear that someone is calling their name. None of these discoveries would constitute mere perceptual knowledge. Each could be accomplished only by inference.” (p. 194)

- 'Perceptual learning' (Gibson 1969) is supposed to be in tension with the tradition that assumes that perceptual experience has unchangeable content. But the experiments only show changes in *cognitive responses* to the sensory data. They do not establish that the changes are accompanied by changes in perceptual experience that are caused by changes in cognition.
- A hypothesis that emerges: "perceptual processing is simply a silent process of translating infosigns that appear in preconscious sensory data into uniceptual representations without uniceptual inference, a process that is fallible and not necessarily accompanied by alterations of sensory experience." (p. 195f.). E.g., as an outfielder, you can see where the ball will land. As a chess master, you can see that black has a mate in five.
- Sometimes, perception proceeds in stages, utilizing signs of signs. But not all translations that could be reconstructed as indirect are in fact indirect.

14.7 Translating Linguistic Signs into Understanding (p.196-197)

- Although adults may have unitrackers and unicepts for letters, words, sentences, and even phonemes, there would seem to be no necessity to call on them in the process of recognizing words to be translated into beliefs.
- Children's resistance to learning through correction illustrates the fact the language learner doesn't hear words as such [but rather what is said].

14.8 Replies to Objections (p.198-203)

14.8.1 But a dog doesn't look like this: "DOG" (p.198-199)

- "Ordinary verbs of perception don't track a brand of psychological processing" (p.198). "'What a dog looks like' suggests a dog directly seen in good like at a short distance" (p.199). But what about a mile away, in a black and white photograph, etc.?
- What does rain sound like when falling on an English speaker? ("It's raining!")

14.8.2 The referents of words like 'electron', 'democracy', and 'intelligence' can be talked about but not perceived (p.199-200)

- "But one who speaks knowledgeably about anything empirically known does so on the basis of evidence that was once carried to some person or some people's senses by infosigns." (p. 199) Nothing is perceived unless through infosigns.
- Consider a CAT scanner: the scanner may allow a technician to see a crack in bone, though the data responsible for the image was gathered over a span of time. Don't confuse the complex and inferential processes involved in producing an infosign with the simple and transparent between the infosign and the infosignified.

14.8.3 Abstract entities can be talked about but can't be perceived (p.200-201)

- Perceiving that something is some shade of red but being unable to tell which [this is not a representation of something concrete, but it is a representation].
- There are no principled restrictions on the range of things that can be perceived.

14.8.4 Understanding without Believing (p.201-202)

- During perception, what is perceived goes directly into belief. But in language, there seems to be a two-stage process: understanding of the utterance, and coming to believe what it says.
- However, experimental evidence suggests that understanding is *not* a separate act taking place before translation into belief. There's a direct line from descriptives to belief and from directives to action.
- Understanding without believing is like rejecting a perceptual illusion. The case of understanding without believing is like that of the man at the door – it takes a retreat to the position that what was perceived was only a spoken sentence [with a specific meaning] [rather than what is the case?]

14.8.5 Referents supplied by speaker intentions must be understood by recognizing speaker intentions (p.202-203)

- Even if there were referents that were fixed by speaker intentions, recognition of those intentions wouldn't be necessary for hearer understanding.
- What fixes/determines the referent of a term is one thing, how the hearer recognizes it is another.

Ch. 15: Markers of Identity and Grounded Infosigns

15.1 Overview

15.2 Selfsigns of Identity: Duplicate Markers; Strawson Markers (p.204-206)

- Knowing what a sign element refers to is knowing which elements in other signs (including both intentional and nonintentional infosigns) carry information about the same thing. One simple way referent identity can be shown is by *duplicate sign elements*: tokens having the same sign design indicate that their referents are the same, that they should be coidentified. "Linguistic signs are specifically designed to allow most same-tracking to be done by attending to repeated design elements." (p. 204)
- In general, whether or not duplication can be used as an indication of cosignification depends on whether the tokens fall within a single reference class that the interpreter is able to recognize.
- Conventional language uses *duplication* to mark identity for individuals, properties, relations, activities, kinds, etc. (Tokens of the sign design "elephant" always refer to the same real kind when the tokens belong to the same REF).
- Additionally, the *identity* of an individual can serve to indicate identity of an individual, kind, etc. (Strawson markers of identity)
- Map example (p. 205): Each addition of a new symbol introduced to show the spatial relation of a new feature to some feature already represented on the map automatically becomes part of a representation of the spatial relation of that new feature to every other feature represented on the map. The new symbol's identity indicates the sameness of the signified feature in each representation of something's relation to it.
- A bee dance says that *the same* nectar source is both x feet away from the hive *and* at 30 degrees clockwise from the sun. (Compare: "Betty is tall and slender" - the

identity of the token of "Betty" indicates the sameness of the person with two properties.)

- Strawson: we should model the store of knowledge one has about a single individual by a single token linked to many predicate tokens at once (like a single dot with spokes radiating out representing the individual's properties, including what it's named).
- A 'Strawson marker' serves to *merge* multiple representations into one.
- An analogue of Strawson's model for the brain's representation of one's known world would present a very different result from the 'language of thought' model. It would obviate imputing to thinkers a multitude of inferences the conclusions of which would already be represented in a Strawson-style brain.

15.3 Anaphoric Signs of Identity (p.206-208)

- An anaphor relates a sign to a cosignifying sign.
- With *asymmetric* anaphors, it is assumed that the identity of one of the signifiers' referents is known independently – "When she woke, Betty went out";
- With *symmetrical anaophors* it is not. "I saw a boy who was selling chestnuts";. here the boy to whom selling chestnuts is imputed is not assumed to be known by the hearer.
- Anaphoric relations between nonintentional infosign elements are always symmetrical. Nature does not know whether an observer can or cannot identify one of it's signs. From here on 'anaphor' will be used to refer to all sign elements related by anaphoric relations (with no distinction beween 'antecedents' and 'postcedents').
- Neatly lined-up hoof prints left by Bambi bear anaphoric relations to one another, showing that it was the same deer that made them. Similarly for recurring tokens of a proper name within a discourse (unless indicated otherwise).
- A chain of weak signs of a red squirrel (falling pine cones, disturbed branches, etc.), occurring in a single space-time line, can allow coidentification and amount (together) to a strong sign. (Likewise, "the rascal" following "the thief" in conversation are likely anaphors).
- The coidentification, or coordinate reference, afforded by anaphoric relations within a conversation is not limited to individual items but can apply to situations, occasions, events, etc. The relations can be sustained by "natural trains of thought of cooperative speakers monitoring their speech, or they may be partly or fully conventional" (p. 208).

15.4 Grounding as Indicating Identity (p.208-209)

- In the case of some infosign elements, recognizing what the element signifies requires recognizing somehow an *'external'* relation between sign and signified – "nothing in the content of the representation … makes it about this object rather than that object" (Dretske 1995). The sign is then read as a *grounded sign* and the state of affairs it signifies is its *ground*.
- "Recognizing the ground of an infosign is interpreting the information that the sign carries as information about that ground as identified in another way, that is, as reidentified." (p. 209). The grounding relation can serve as another kind of sameness marker if the sign reader understands what it is.

- Proposal: incomplete descriptions and quantified phrases that have seemed puzzling are not intentional signs of their extensions, but rather infosigns whose function is to indicate that their tokens should be read as grounded signs. (To be unpacked in Ch. 16.)

15.5 Situated Signs; Counting up Signs (p.209-210)

- Grounded signs (and many anaphoric signs) are *situated* signs: they require considering the situation in which the sign is produced.
- Anaphoric relations within a discourse can hold between whole sentence tokens (not just within them). Whatever elements have been purposefully produced in relations to one another so as to be understood together count as being a single sign. A lesson: 'sign' has no clear individuation conditions.

15.6 Recognizing Identity: A Reminder (p.210-211)

- Earlier (3.4) it was claimed that "all identifying is really coidentifying". Using the example of grounded signs, it can now be appreciated that, unless one is able to *re*identify e.g. the deer who made the hoof print, one cannot even begin to have *knowledge which* deer that was. This is contra Russell and Evans, who imply that it would be sufficient simply to identify it via the description *the deer who made this print.* A definite description gives a way that the thing could be identified but if it never gets *re*identified, which one it is remains unknown.
- "You begin to know what a thing is only when you begin to recognize other sign tokens of it that fill in more information about it. … Knowing what a thing is is not a mere ability, but the fruit of having exercised that ability." (p. 211).

Ch. 16: Out-Side Pragmatics: Descriptions, Quantifiers, Directives

16.1 Overview

16.2 Three Kinds of Referent (p.212-213)

- *Outside pragmatics*: concerns cases where the referent of a sign element lies outside the semantic content of the utterance, so the hearer must identify the sign's ground to interpret it. How does the hearer do that – e.g., tell what "Peter" refers to?
- Three kinds of content: speaker content, intentional (=semantic) content, and natural informational content (13.4) ? 3 kinds of referents.
- *Hybrid constructions* (e.g. "Jane led **everyone** to **the park**") ("I'd like to read **Peter's book**") do not have semantic referents. BUT their referents are *not* supplied by speakers either. Rather, they are conventional devices whose referents are determined by their natural *informational* content. Compare knowing what has the voltage when reading a volt meter
- The three kinds of reference can come apart – e.g., in abNormal cases where an utterance is false (no natural referent), or when a speaker uses a wrong word/name (speaker content from semantic content). But they also come apart in some Normal cases, where the natural referent matches the speaker's but there's no semantic referent at all (hybrid constructions). ("Where did **the boys** go?")

- In such (hybrid) Normal cases, the referent is found in the ground of the sentence in which the construction occurs. The construction's intentional content provides only a guide to the ground that the hearer must look for (outside the sentence).

16.3 Functions of the Definite and Indefinite Articles (p.213-216)

- The indefinite article "a" and the definite article "the" do not, of course, have referents/extensions. Their contribution to the semantics of sentences in which they occur needs to be understood in terms of their *stabilizing functions*: "an effect on the compliant hearer that advances communication with frequency enough to account for continued speaker use and hearer compliance" (p. 213f.). Articles of this kind are found only in some languages. Their functions are not mandatory, but it is obviously useful to have them.
- A pure guess that Triumph will win a race *that (accidentally) turns out right* has a truth maker – the state of affairs that would have had to be the utterance's ground if it had had one. But a pure guess has no ground. [A truth maker is tied only to the semantic rules for the utterance.]
- *Contrast*: "I met a man from Utrecht yesterday". Here the semantic rules do not determine a referent for "a man". But "[i]f this utterance were Normally true (true in accordance with a Normal explanation for arriving at truths) a natural referent would be supplied." (p. 214). To see this, consider a case where, unknown to me, I met two men from Utrecht yesterday. Only 'the man I had in mind' would be the ground, and thus the natural referent of my token of "a man", since he is what my utterance concerns *as an infosign*.
- *Proposal*: "the stabilizing function of the indefinite 'a' in a description … is to signal to the hearer that he need not identify the (presumed) natural referent, the *ground* for the description token" (p. 214). Recovering the relevant information in an utterance containing "a" does not require recovering its ground.
- The stabilizing function of "the", by contrast, is to signal to the hearer that a referent does need to be identified for the description token. If no antecedent is supplied for "the", it is signaled (semantically) to the hearer that she is expected to retrieve the description's ground. This can be accomplished by finding a target of joint attention, a focus of previous discourse, what the speaker has been attending to, etc.
- **"But** however we find this out, it need not be by thinking of the speaker's intentions. Any mark that correlates well with where certain kinds of information originate will do" (p. 235).
- If I say "I met the man from Utrecht today" where there was no such man grounding my utterance, what I said lacks a truth condition. (Also, of course, a ground.)
- A definite description that is not unique is a hybrid of an intentional sign and a nonintentional infosign.
- Donnellan's distinction between 'referential' and 'attributive' uses may amount to "the distinction between those uses where the speaker's purpose requires that the hearer identify something already known about, and those uses where the description, being unique, merely *could* be used for reidentification" (p. 216).
- The use of a definite description that is not obviously unique is usually an infosign that the natural referent is either very salient or else unique under its description.

- Re Moore's paradox ("It's raining but I don't believe that it's raining"): "The act of uttering a descriptive sentence is usually an infosign that one believes it, so uttering a sentence, 'p', tends to absorb as part of its *intentional* content that I believe that p." So the conjunction is "a halfway contradiction". (p. 216)
- Here, as with use of "the" with a description that isn't obviously unique and where the natural referent is not salient, we see the 'interplay between the semantic content of a linguistic form and the natural information carried by the act of using that form" (see footnote 2).

16.4 Quantifiers, Possessives, and Proper Names (p.217-218)

- When quantifiers are used along with "the", hearers have to figure out (and it's generally very obvious) whether a restricted domain needs to be identified as the natural ground for the utterance.
- Again, the hearer of "Peter's book" has to recognize the possessive's ground □ *book bearing the pairing relation to Peter,* prompting the hearer to identify the more specific relation that is being talked about, though it is not semantically identified.
- Again, no need to think about the speaker's mind to retrieve the relevant ground. [*Consider: I can use your gaze to learn which object you want me to retrieve – as do dogs – without thinking about your gaze or* reflecting on the object that you have *in mind*, as such.]
- When it comes to proper names, again, semantic referent is not *determined by* the referent the speaker has in mind. A token of "Bill" will have as its semantic referent the referent of the particular REF from which it has been copied.
- "Understanding where another's mind is focused can be rather like understanding where a telescope or a television is focused, allowing one to interpret where the information being dispensed originates but without understanding or thinking about the intervening process. It requires no more understanding of what is inside a mind than the child knows of what is inside the camera. Humans are very good at following the focus of another's mind to understand their speech as well as their other activities" (p.218).

16.5 Thumbnail Review of Basic Themes (p.218-220)

[This section should be read in full!]

The themes:

- Real [worldly] kinds as the starting point
- Extensional terms as directly referential
- Unitrackers as nature's solution to the most fundamental challenge for cognition (recognizing sameness)
- Unicepts (affording, substantive, attributive) as unshared repositories of collected knowledge
- AbNormal circumstances as resulting in improper functioning of unicepts (false beliefs, equivocepts, redundant unicepts, vacucepts)
- Infosigns and intentional signs as overlapping kinds of signs
- Understanding of linguistic signs as subsumed under understanding of natural sign

- Intentional content as semantic meaning, and not dependent on speakers' intentions but on the REF from which it comes
- Natural language interpretation as of the same kind as perception (which is translation into mental representation)