What Model Theory Does, and Does Not Do

ESSLLI 2014

Michael Glanzberg

Northwestern University

Background Issue

- Model-theoretic semantics.
  - Highly successful. Maybe the mainstream approach.
  - Most obvious success in what is sometimes called 'compositional' semantics.
  - Will ask what the role of model theory in model-theoretic semantics really is.
- Lexical semantics.
  - Huge range of approaches and techniques.
  - A little model theory (e.g. Dowty, 1979).
  - But more often, representations of meaning more like what we find in cognitive psychology.
  - Or, computer science, . . .
- General background issue: should these be integrated? How?

Goals for Today

1. Discuss the place of model theory in semantics.
   - Reconsider where model-theoretic techniques fit into semantics, and contrast with the place of disquotation.
   - Distinguish: Mathematical semantics from model theory proper.
   - Mathematical semantics can be useful for providing interesting generalizations and explanations.
   - Both type-theoretic and neo-Davidsonian approaches can and do make use of mathematical semantics, but also rely on disquotation at key points.
2. Consider some ways that truth-conditional approaches fail to capture aspects of lexical meaning.
3. Introduce a view of the lexicon combining truth-conditional and psychological explanations of lexical meaning.
4. Explore how they interact.

A hybrid of old work (forthcoming; 2011) with work in progress.

The Plan

Model Theory, Truth Conditions, and Semantics

Mathematics and Disquotation

What's Left Out?

Concepts and the Lexicon

A Little History

- Logic and natural language in the 60s and 70s.
  - Montague's (e.g. Montague, 1973) idea that model-theoretic techniques could be applied to the study of natural language semantics.
  - A competing idea from Davidson (1967). Apply techniques from Tarski (1935) to natural language semantics.
- Towards today.
  - Both projects have been taken up by a number of linguists and philosophers.
  - To some, logic and language are a natural combination, e.g. the textbook of Gamut (1991).
  - Both Montague's and Davidson's proposals have taken root in semantics, e.g. the textbooks of Heim & Kratzer (1998) and Larson & Segal (1995).
  - Leading to . . .

Two Approaches to Truth-Conditional Semantics

- Widely held idea that within the broad area of truth conditional semantics there are two distinct approaches.
  - **Absolute semantics** (Neo-Davidsonian Semantics). Derives clauses like: 'Ernie is happy' is true $\longleftrightarrow$ Ernie is happy.
  - **Model-theoretic semantics** (Montagovian semantics). Derives clauses like: For any model $\mathfrak{M}$, 'Ernie is happy' is true in $\mathfrak{M} \longleftrightarrow \mathit{Ernie}^{\mathfrak{M}} \in \mathit{happy}^{\mathfrak{M}}$.
- Both are truth-conditional, not, e.g. conceptual role, cognitive, etc.
- Commonly held that even so, they are very different sorts of theories.
- Has been argued by Lepore (1983) that model-theoretic semantics is somehow defective, or at least less satisfactory than absolute semantics.
- See also Cresswell (1978), Higginbotham (1988), and Zimmermann (1999).

Starting Points I

To begin, I shall discuss Lepore's argument. He makes a few opening claims that I shall accept without argument.

- Semantics must provide an account of what a speaker understands about their sentences, i.e. what they know when they know what their sentences mean.
- Take this to be part of competence: the main subject-matter of an enterprise in linguistics.
- Semantic competence is a species of this.
- Will occasionally talk about competence in very Chomskian terms. Will ultimately need at least domain-specificity.

Starting Points II

- Truth conditions provide a central aspect of semantic competence.
- NB will revisit the strength of this conclusion late, but assume for now.
- *Disquotational statements* state truth conditions, and so offer non-trivial statements of key components of speakers' linguistic competence.

Lepore's Argument I

Lepore argues:

- Model-theoretic semantics does not provide truth conditions in any way that accounts for such understanding.
    - Only provides *relative* truth conditions: i.e. conditions for truth in a model.
    - Much weaker than disquotational facts.
    - Fails to really provide truth conditions at all.
    - Thus fails to explain what speakers know about the meanings of their words and sentences. Fails to explain competence.
- The problem with relative truth conditions is it leaves open what our words refer to or mean, and so, fails to really fix the truth conditional knowledge that is part of our linguistic competence.

Lepore's Argument II

- We get the right truth conditions if we provide absolute statements of reference and satisfaction properties, i.e. 'Ernie' refers to Ernie and 'is happy' is satisfied by something iff it is happy.
- Thus, we get absolute truth conditions if we provide an absolute semantics.
- Conclude: absolute semantics succeeds, model-theoretic semantics fails.
- Why not take a more model-theoretic path, and provide absolute truth conditions by specifying which model is the correct or 'intended' one?
    - This is too demanding: it requires huge amounts of factual information, beyond what any speaker could be expected to know.
    - It also requires knowledge of facts about complex mathematical objects, like functions, intensions, etc., which speakers do not know.

Official Reply

- Current 'model-theoretic' semantics is really *absolute*.
- No specific reference to models in 'model-theoretic semantics', e.g. not in Heim & Kratzer (1998) or Chierchia & McConnell-Ginet (1990).
- What we do find in the textbooks is something like:
    - $[\![\mathrm{Ann}]\!]$ = Ann, $[\![\mathrm{smokes}]\!] = \lambda x \in D_e.\, x$ smokes.
    - Function application: If α is a branching node, β, γ its daughters, then $[\![\alpha]\!] = [\![\beta]\!]([\![\gamma]\!])$ or vice-versa (Heim & Kratzer, 1998; Klein & Sag, 1985).
- Does state 'absolute' facts about truth and reference, not relative to a model.
- Actually, older presentations which officially relied on a notion of truth in a model (e.g. Dowty et al., 1981) usually drops reference to it when the linguistic analysis starts to get interesting.
- So, not really a target for Lepore.

What Is Distinctive about the 'Model-Theoretic' Approach? I

- What seems distinctive about this approach?
  - Use of the λ-calculus.
  - Assignments of semantic values to all constituents.
- Assigning semantic values is not (I claim) the semantically important feature. No important difference between saying:
  - $[\![\text{smokes}]\!] = \lambda x \in D_e.\, x \text{ smokes}.$
  - $Val(x, \text{smokes}) \longleftrightarrow x \text{ smokes}$ (e.g. Larson & Segal, 1995).

  Fact that there is an object in the metatheory for the first does not change the way it describes contribution to truth conditions that speakers know.

What Is Distinctive about the 'Model-Theoretic' Approach? II

- Even so, the use of λ's is significant.
    - A theory of semantic *composition* in terms of functions and arguments.
    - Some sorts of semantic analyses become available, such as higher-type modifiers (e.g. *each* as a VP modifier , not a quantifier).
    - Potentially far-reaching implications for logical form (cf. Pietroski, 2005).
    - Different use of events (and theta roles) in analyses.
- In fact, neo-Davidsonian and 'neo-Montagovian' theories often provide different analyses and make different predictions.
- But they need not differ on the fundamental project of truth-conditional semantics.

Where Are We?

- Conclusions so far:
  - There are no fundamental differences between 'model-theoretic' (type-theoretic) and 'absolute' (neo-Davidsonian) approaches to semantics in how they give truth conditions.
  - There are empirical differences, surrounding use of higher-types, etc., and their relations to meaning and LF.
- Where next?
  - Consider where our semantic theories provide explanations.
  - Argue that disquotation fails to offer needed explanations, where model theory and other applications of mathematics succeeds.

Where Next?

Model Theory, Truth Conditions, and Semantics

Mathematics and Disquotation

What's Left Out?

Concepts and the Lexicon

What Is Bad about Disquotation I

- Even though they are non-trivial, disquotational statements are *boring*!
  - They can be generated with only minimal knowledge of the semantic category into which an expression falls (cf. Higginbotham, 1989a).
  - They state semantic facts in a way that fails to offer interesting generalizations, or other aspects of good *explanation*.
  - In particular, will not explain anything about the meaning of any expression that was not already transparent to you.
- Return to the theme of Field (1972):
  - Mere lists of facts fail to yield good explanations.
  - Disquotation merely lists substantial facts about truth conditions.
  - As Field pointed out, list is trivial to generate, and fails to explain anything more than saying that *S* has the truth conditions it does.

What Is Bad about Disquotation II

- Use of higher-type objects does not change this.
  - Each of these is equally weak in explanatory power:
    - $[\![\text{smokes}]\!] = \lambda x \in D_e.\, x \text{ smokes}.$
    - $Val(x, \text{smokes}) \longleftrightarrow x \text{ smokes}$
  - Each fails to be able to explain anything about the meaning of *smokes* that was not already transparent to you.
- Moral: either type-theoretic or neo-Davidsonian theories run the risk of being true but explanatorily vacuous theories of linguistic competence.
- The real worry from Lepore's argument: we need something like disquotation to genuinely state truth conditions, but most of the ways we know how to do that threaten to be explanatorily vacuous.

Where Model Theory Actually Does Something

- Higginbotham (1988): model theory as the lexicography of the logical constants.
- Common example: the semantics of determiners (Barwise & Cooper, 1981; Higginbotham & May, 1981; Keenan & Stavi, 1986).
    - $[\![\mathrm{most}]\!](A,B) \longleftrightarrow |A \cap B| > |A \setminus B|$.

  Determiner denotations are sets of sets, or corresponding functions.
- A theory of determiner denotations that makes non-trivial predictions and generalizations. For instance:
    - The conservativity universal for determiners.
    - The Ladusaw-Fauconnier generalization.
- These could not be stated for a purely disquotational account of determiners, though can be done in either Montagovian or neo-Davidsonian frameworks.

But Is it Really *Model*-Theoretic?

- Though GQ theory is model-theoretic in that it uses the tools and techniques of abstract model theory, it does not always assume that the model itself varies.
- Need to do some elementary set theory, but can do it with real-world objects.
- Hence, avoids the Lepore objection.
- Not using model theory as a framework.
    - Applying some mathematics to problems in semantics.
    - Right mathematics helps build theories with clear explanatory pay-offs.
    - An instance of the 'unreasonable effectiveness of mathematics' (Wigner).
- Use of mathematics produces an explanatorily rich theory of the semantics of quantifiers, and other logical expressions.

Local and Global I

- The theory of generalized quantifiers captures the role of *model theory proper* in the difference between *local* and *global* generalized quantifiers.
  - Local: $[\![\mathrm{most}]\!]_M = \{\langle A, B\rangle \subseteq M^2 : |A \cap B| > |A \setminus B|\}$.
  - Global: function from M to $[\![\mathrm{most}]\!]_M$.
- For studying model theory proper (i.e. the study of abstract model theory as a branch of logic), global GQs are the basic notion.
- For studying semantics, the basic view is local, with the domain taken to be whatever is in the real world.
- For some semantically relevant results, we must take a local perspective (e.g. Keenan-style results, like that over a finite universe, every type $\langle 1\rangle$ GQ is denotable by a complex English Determiner).

Local and Global II

- In some cases, local results are available, but global ones are stronger (e.g. Westerståhl, 1989):
  - CONSERV (local): for every $A, B \subseteq M$, $\mathbf{Q}_M(A, B) \longleftrightarrow \mathbf{Q}_M(A, B \cap A)$.
  - Weaker than a genuinely global UNIV: for each M and $A, B \subseteq M$, $\mathbf{Q}_M(A, B) \longleftrightarrow \mathbf{Q}_A(A, A \cap B)$.
- Hence, it may sometimes be useful to appeal to model theory proper as part of our mathematical semantics.
  - Model theory is used as just another mathematical tool for mathematical semantics.
  - Does not change the basic approach to truth conditions.
  - Useful in other settings, e.g. as a tool to study entailment, etc.

Morals from the Lexicography of Logical Constants

- In some cases, including the logical terms, use of mathematics provides explanatorily substantial theories where disquotation does not.
- Model theory can provide tools for doing so, without abandoning absolute truth conditions.
- As before: type-theoretic and neo-Davidsonian theories can both use mathematics in semantics, though may in some cases use it differently.

Beyond Logical Terms

- Mathematics (often relying on type theory) has been applied widely in the 'compositional' side of semantics.
- Also to various operator-like expressions or functional ones (e.g. tense and mood, adverbs, focus, etc.).
- What about the heart of lexical semantics: verbs, nouns, adjectives?
- Extending Higginbotham's view, I shall argue that we do find some applications of mathematics to the semantics of these expressions as well.
- However, not as fully analyzable in mathematical terms as the logical expressions.
- Present the example of gradable adjectives.

Gradable Adjectives I

Present an analysis of gradable adjectives that allows for some interesting mathematical explanations (Barker, 2002; Bartsch & Vennemann, 1972; Cresswell, 1977; Kennedy, 1997, 2007; von Stechow, 1984).

- Core meaning of an adjective is a function from individuals to degrees on a scale.
    - $\mathbf{tall}(x) = d$ a degree of tallness.
    - Scale: an ordered collection of degrees, with a dimension specifying what the degrees represent.
- Scale structure is useful compositionally:
    - Comparatives and positive morphology.
    - Measure phrases and degree terms like *very*.

Gradable Adjectives II

- Relevance to lexical semantics:
  - Invertability of ordering and antonym pairs like *short/tall*.
  - Dimension can help explain which adjectives can combine in complex comparisons. Suppose *wide* and *tall* both have scales of linear extent, while *flexible* does not (Kennedy, 1997; Kennedy & McNally, 2005):

    (1) a. He is as wide as he is tall.

    b. ?? He is as tall as he is flexible.

- A little bit of mathematics can be applied to get further explanations (Kennedy & McNally, 2005).
  - Scales can be closed or open.
  - Explains differences in meaning between absolute adjectives like *closed* and relative ones like *tall*.
    - Entailments: *The door is not open* entails *The door is closed*, but *The door is not large* does not entail *The door is small*.
    - Proportional modifiers: *half closed*, **half tall*.

Gradable Adjectives III

- Mathematics can help formulate substantial empirical generalizations.
- But, we still have a *nearly* disquotational component: the dimension.
  - The dimension for *flexible* is degree of flexibility or bendability.
  - The dimension for *tall* is tallness or linear extent along the vertical axis. (How far from disquotation?)
- With the logical constants, mathematics provided an essentially complete semantics theory.
- With genuine lexical elements, it is one part the semantics.
  - Does not replace stating referential properties in languages already endowed with those sorts of properties.
  - Can supplement and refine those statements, and build generalizations around them.

Conclusions So Far:

- Both current 'model-theoretic' (type-theoretic) and neo-Davidsonian semantics state absolute truth conditions.
- Both can fail to build explanatorily substantial theories when they fall back on disquotation.
- The application of mathematical tools and techniques—is one way that both type-theoretic and neo-Davidsonian theories can formulate more explanatorily significant claims.
- These techniques are often drawn from model theory, but not used as they are in model theory proper.

A Loose End

- Always mathematics?
- What about less mathematical 'elucidations' of features of meaning.
- Example: the famous characterization of the meaning of *cut* as a linear separation of the material integrity of something by an agent using an instrument (Hale & Keyser, 1987; Higginbotham, 1989b).
- Seems to be of some explanatory value.
- Would a more refined theory couched in mathematics be better?
- Not clear. But at least, mathematics one powerful way to get explanations.

Where Next?

- Examine what the kinds of good, typically mathematical, explanations we have seen can provide.
- Consider what happens when we fall back on disquotation, and what we do not account for when we do.
- Use that to motivate an interface view of the lexicon.

Where Next?

Model Theory, Truth Conditions, and Semantics

Mathematics and Disquotation

What's Left Out?

Concepts and the Lexicon

Mathematics and Semantics I

- What mathematics does in an empirical semantic theory.
    - Allows us to formulate descriptions of specific, typically more abstract, features of lexical or compositional meaning.
    - These capture aspects of speakers' competence. Usually highly tacit, as not readily transparent to many speakers.
    - Features usually give only partial accounts of contributions to truth conditions (e.g. scale structure, etc.).
    - Quantifiers (and other logical terms) are unusual in that we can give essentially complete accounts of contributions to truth conditions.
- Goal is partial description with explanatory value.

Mathematics and Semantics II

- Mathematics is the tool for doing so.
  - In this respect, not directly significant whether we use types, scales, Boolean algebras, sets, or meaning postulates, etc.
  - But can acquire significance if has empirical consequences (e.g. types and LF), or allows or impedes formulating generalizations (e.g. role of scales).
- Disquotation marks places where explanations in our theory give out.

Partiality

- Where our theories rely on disquotation, they fail to offer good explanations.
- With lexical categories, our theories always fall back on disquotation (or near disquotation) at some points.
- Logical constants are an exception.
- Thus, in general, our semantic theories are explanatorily partial.
- Conclusion I defend elsewhere (Glanzberg forthcoming):
  - This is because our semantic competence really only encodes partial contributions to truth conditions.
  - I.e. partiality in competence!

What Is Left Out? I

- Our theories give out at key places for the lexical categories.
- Leave out the core meanings of sentences in terms of objects, properties, and the way the latter apply to the former.
  - Core explanations of specific properties corresponding to predicates.
  - Core explanations of kinds of objects corresponding to terms.
- Seems to leave out the 'core conceptual meaning' of our terms.
- For these, we get disquotation, or near disquotation, in our semantic theories.

What Is Left Out? II

- Need to find some way to explain what disquotation is doing in our theories, and make room for other sorts of explanations of core meaning.
- Maybe draw on other approaches to lexical semantics or psychology?

Two Types of Meaning?

- Areas where there have been substantial explanations drawing on mathematical semantics include:
  - Functional elements.
  - Structural elements within lexical items.
- Structural-functional: yields to substantial linguistic characterization.
- Core conceptual: seems not to.
- Why two types?
  - Conjecture that only the structural part is in the language faculty proper, and hence, the only place where linguistically driven mathematical analyses pay off.
  - A form of partiality in semantic competence.
- Need to fit the two types together.

An Interface Picture I

- A view of the lexicon that can make sense of the two kinds of meaning and the places we find successful explanations in semantics.
- An interface view: are disquotational statements are marking points where the language faculty proper calls wider conceptual resources.
  - Core concepts not language-faculty specific.
  - Meanings built by combining core concepts with language-faculty-specific components.
  - Latter are described by mathematical semantics and related theories.
  - Yield substantial linguistic generalizations.
  - Disquotation statements mark points where such core concepts are introduced.
  - Do not yield substantial linguistic generalizations.

An Interface Picture II

- Pointers and packaging.
  - The language faculty has interface points to wider conceptual faculties.
  - (For externalists, may extend outside the individual.)
  - Lexical meaning includes these.
  - Disquotation enters at those points?
  - Provides *pointers* to extra-linguistic concepts.
  - These are *packaged* by structural-functional elements, within the language faculty.
  - Pointers point to concepts which function as lexical roots.
  - For linguistic theory, function as atomic (cf. Grimshaw, 2005).

Implementing the Interface Picture I

- Lexical roots are treated as *pointers* to concepts.
- These are *packaged* by lexical entries.(Glanzberg, 2011, forthcoming) (cf. Pietroski, 2012, 2010).
  - Give each lexical item its distinctive, idiosyncratic meaning.
  - Give it rich conceptual structure.
- Packaging can be implemented by the rich tradition in lexical semantics (especially for verbs) that lexical entries are structured (Bierwisch & Schreuder, 1992; Pinker, 1989; Jackendoff, 1990; Wunderlich, 1997)
- An influential version from Levin and Rappaport Hovav (Levin & Rappaport Hovav, 1995, 2005; Rappaport Hovav & Levin, 1998).
  - An event decomposition approach.
  - Predicate decomposition within the lexical entry describes decomposition of an event into structural components.

Implementing the Interface Picture II

- Example, *open*: 'externally caused change of state'.

(2) a. open

b. [[*x* ACT] CAUSE [BECOME [*y* ⟨*OPEN*⟩]]]

- Features of the analysis: packaging within the lexicon.
  - A root element *OPEN*.
  - An event-structural frame, built from elements including CAUSE and BECOME.
  - Limited number of structural elements.
- Grammatically relevant facts of meaning.
  - Structural frame predicts aspects of grammar: especially, argument projection and argument alternations for verbs.
  - Explain how lexical items group into classes.
  - Can also capture some entailments.
  - Packaging is thus taken to be distinctively linguistic.
- Other approaches with similar goals, notably Hale & Keyser (1993, 2002); Harley (2007).

Where Next?

- We now have the pointers and packaging view.
- Allows a domain for fruitful, typically mathematical explanation in semantics.
- Includes pointers to extra-linguistic concepts.
- We should now consider what the targets of those pointers can do for our understanding of meaning

Where Next?

Concepts and Roots

- Roots are linguistically atomic.
- Treated as disquotational by our semantic theories.
- But, are pointers to *concepts*.
  - Provide core conceptual meaning for lexical items.
  - Provide the distintive idiosyncratic meaning of the item.

- Concepts are:
  - Mental representations.
  - Units active in thinking.
  - Contribute content to cognitive states.
  - Part of our broader cognitive repertoire.
  - Some concepts appear to be acquired early enough to be pre-linguistic.
- Assume concepts (typically) characterize real-world things.
  - Concepts determine *categories* of things.
  - Complex issues about how to divide up the referential and internal aspects of concepts (e.g. long-arm versus two-factor theories).

Conceptual Structures I

- Current research on concepts, mostly from cognitive science, considers a range of different kinds of conceptual representations, each of which offers a view of the nature of concepts. (See surveys by Laurence & Margolis (1999) and Murphy (2002).)
- Prototype views, stemming from work of Rosch (e.g. Hampton, 1993; Prinz, 2002; Rosch, 1978; Rosch & Mervis, 1975; Smith & Medin, 1981).
-
    - Features found in categories.
    - For BIRD:

fly	*wings*
feathers	*lay eggs*

Conceptual Structures II

- Features are *weighted*. Empirical results about *typicality* effects.
- Categorization is done by applying some *similarity metric* that compares weighted features.
- Gives concepts a *graded* or *probabilistic* nature.
- Glossing over a lot of issues about how these are represented (exemplars versus prototypes, etc.).

- Theory views, stemming from work of Carey (e.g. Carey, 1985, 2009; Gopnik & Meltzzoff, 1997; Keil, 1989; Murphy & Medin, 1985; Rips, 1989).
  - Concepts are mental theories.
  - Relate multiple concepts.
  - Involve laws and explanatory mechanisms.
- Of course, many other options, and many complications.
- Enough to give us a sense of what concepts might be like from a psychological perspective.

Words and Concepts I

- Rich conceptual structure seems to be the right sort of thing to more fully articulate the meaning of a lexical item.
- Offer rich theories of core conceptual meaning.
- Also, from cognitive science, so evidence connecting words and concepts.
- From the survey by Vigliocco & Vinson (2007, p. 195), “word meaning must be grounded in conceptual knowledge.”
    - Connections between speaking and thinking (e.g. Murphy, 2002).
    - Priming and typicality effects (e.g. Murphy, 2002).
    - Imaging data on activation suggesting: “When we activate semantic representations we also activate conceptual information” (Vigliocco & Vinson, 2007, p. 196).
    - Word learning associates linguistic forms with concepts (Bloom, 2000; Clark, 1983).

Words and Concepts II

- But, difficult to decide how direct the mapping should be.
  - Semantic and cognitive deficits come apart (Vigliocco & Vinson, 2007).
  - Complex issues about cross-linguistic lexicalization differences. Highly controversial. (Reviewed in Bowerman & Levinson (2001), Gentner & Goldin-Meadow (2003), Bloom (2000), Carey (2009), Pinker (1994).)
- The interface picture offers a way of establishing a complex mapping between lexical meanings and concepts.
- Offers a way to combine structural-functional and core conceptual aspects of meaning.
- Provides separate domains of explanation for each.
- Specifically, can distinguish distinctively linguistic from broader cognitive aspects of meaning.

An Example I

- Look at the old stand-by *kill*.
- Consider issues that arise as we work through the pointers and packaging idea.
- Lexical entry, in Levin and Rappaport Hovav style:

(3) a. $[[x \text{ ACT}] \text{ CAUSE } [\text{BECOME } [y \ \langle \mathit{STATE} \rangle]]]$
    b. $[[x \text{ ACT}] \text{ CAUSE } [\text{BECOME } [y \ \langle \mathit{DEAD} \rangle]]]$

  - *STATE* marks the position of the root.
  - For *kill*, get root *DEAD*.

- Compositionally, root needs to act like a predicate.
- The problem we face is that semantic composition requires a predicate P, but what we have is a pointer to a concept that provides a *STATE*.
- How can a concept give us the right predicate?

Answers from the Pointers Approach I

- Not really a problem for the pointers view, but does have some consequences.
- Pointers need not, and generally will not, reflect the internal structure by which a concept is represented cognitively.
- For purposes of semantic composition, they are variables of appropriate semantic type.
- Actually, in the daily practice of doing compositional semantics, this is how roots are treated
- The pointers view holds that the practice is in fact correct.
- The concept to which a pointer points constrains the value of the variable.
    - Structure of the concept does not enter into semantic composition directly.
    - Only via value of a variable.
    - Structure of the concept will provide constraints on that value.

Answers from the Pointers Approach II

- Back to our example.
  - Our root of a causative verb calls for a *STATE*, e.g *DEAD*.
  - Regardless of internal structure of concept DEAD, have a variable S.
  - Compose as $S(x)(e)$
  - Pointer tells us value of S constrained by $S(x)$ iff x falls under *STATE* of DEAD.
  - I.e. $S(x)(e)$ iff x is in the category fixed by DEAD (in e).
- MORAL: once we have the pointers conception, the specific problem about concepts figuring into semantic composition goes away.
- But we will see, that does not mean all problems go away.

Concepts, Categories, and Variables I

- How do concepts constrain variable values?
  - Already noted, concepts determine categories of real-world objects.
  - These are extensions.
  - Pointer stipulates mapping of variable value to this extension.
- But, several issues remain. To illustrate, look at our example of a root concept, DEATH.

Concepts, Categories, and Variables II

- Chosen because a great deal is know about this concept from the psychology literature, from Piaget onward. I follow Carey (1985); Slaughter et al. (1999).
  - Young children represent some forms of biological concepts.
  - Mature concepts of life and death emerge by around age 10.
  - Argued to be represented as a *theory*.
  - Components include:

(4) a. Applies only to living things.
b. Irreversibility.
c. Cessation of bodily function. Biological.
d. Inevitability. Part of life cycle.
e. Caused by breakdown of bodily function.

Concepts, Categories, and Variables III

- Some developmental probes.
  - Younger children treat death as a being in some other place or state, and so not fully incorporating irreversibility.
  - Younger children have different biological concepts, liked to behavior. So, difficulty applying the concept to plants. Missing Cessation (e.g. Nguyen & Gelman, 2002).

Concepts, Categories, and Variables IV

- Back to pointers.
  - DEATH appears to be represented by a theory, incorporating generalizations like (4).
  - Not the kind of structure that can easily figure into semantic composition.
  - We need a predicate $P(x)(e)$ with an extension, not a theory.
  - Root $DEAD(y)(e)$ is variable P, with stipulation $P(x)(e)$ iff x falls under DEAD in e.

Concepts, Categories, and Variables V

- In this case, the concept appears well-suited to this task.
  - Pointer simply tells us variable's value is to be all and only the things that fall under the concept.
  - Example: theories constrain what falls under them, by telling us what those things have to be like.
  - x is dead iff x was living and has irreversibly ceased bodily function, etc.
  - This appears to fix a category which is appropriate to be an extension.
  - This is the value of the variable.

Concepts, Categories, and Variables VI

- Does not really matter if we invoke structure of theories.
  - In this case, thinking of the concept as a collection of features fixing a prototype appears to do equally well.
  - The category is simply those things that have the features, which is an adequate extension.
- General moral: no need to put the internal structure of the concept—the way it is represented cognitively—in the lexicon.
- That is the payoff of the pointers view.

Where Issues Remain: Extensions I

Will most concepts determine (categories that provide) extensions?

- NB we only need something weaker: the concepts that figure as roots, e.g. monadic stative concepts for many verbs, need to.
- Seems to work for concepts represented by reasonable intuitive theories, like DEAD on the Carey story.
- But issue remains.
  - What about concepts that show stronger gradation or typicality effects.
  - E.g. BIRD, which shows gradation on cases like penguins.

Where Issues Remain: Extensions II

- Options:
  - Find a way to fix extension in context, e.g. fix a particular similarity metric and cut-off point?
  - Work with a more complex semantic value that is able to show effects of gradation, e.g. many-valued predicate or something with scale structure (cf. Kamp & Partee, 1995)?
  - Appeal to other forms of contextual input?
  - Work with locally constructed concepts (in extremis, Barsalou (2003)?
  - View compositional semantics as only a partial determinant of truth conditions (cf. Pietroski, 2003)?

Where Issues Remain: Extensions III

- In lieu of an answer:
  - The general issue here is a deep and general one: how much our thinking or speaking really is able to divide up the world into sharp categories.
  - I would not expect lexical semantics to answer this fully. We need more on how concepts relate to the world to get a good answer.
- My own preference is the first two options.
  - I build in a form of partiality: the linguistically encoded aspects of a lexical entry are only partial determinants of truth conditions.
  - But, I would like the combination of a pointer and its packaging to be a full determinant.
  - General reasons that truth conditions seem to be the right kinds of contents for what we say, allows for modeling entailment, etc.

Where Issues Remain: Externalism I

Issues about externalism.

- My proposal about *DEAD* makes its extension internally determined, by a theory a speaker represents.
- This seems right (to me) for the verb *kill*.
- But, not clear we can always be so internalist.
- Familiar examples.
  - Social deference for technical terms. Speaker will not represent a theory or prototype, etc.
  - Content externalism for kind terms. Even if speakers internally represent a theory for FISH that makes wales fishes, they are not, and should not fall in the extension.

Where Issues Remain: Externalism II

- Option: Might be features of conceptual representations.
  - For instance, maybe concepts whose representations make room for substantial psychological essentialism map to categories that depart from the theories represented.
  - Might have partial knowledge of a concept invoke social deference, e.g. ‘whatever the experts say’ (cf. Higginbotham, 1989b).
- Option: Might be lexically determined in some cases.
  - Might be that lexical entries for e.g. natural kind terms bypass concepts for fixing the vales of the relevant variables.
  - Might mark technical terms for partial knowledge?
  - Would make for significant lexical differences between terms like *kill* and like *fish*.
- Potential for different degrees of externalism, either for different concepts or different lexical entries.

Where Issues Remain: Externalism III

- These are all hard and important problems. The packaging view does not solve them.
- But, it still has some virtues
  - By putting aside the problem of concepts composing truth-conditionally, allows us to focus on these problems about the nature of content and linguistic meaning.
  - Thus, puts the problems where they really belong.
  - Allows some flexibility. We can formulate a rich range of theses.
  - Allows for a range of options. Different ways of balancing lexical and conceptual contributions may apply to different cases.

Conclusions

- A role for model theory, as an instance of applied mathematics in semantics.
- Pointers and packaging:
  - Roots are pointers to concepts.
  - Semantically provide variables.
  - Concepts constrain the values of those variables.
  - Variables figure in linguistically determined frameworks that package them into lexical entries.
- Explanations:
  - Standard mathematical explanations in semantics, within linguistics, explain packaging.
  - Psychological explanations explain the targets of roots.
- Problems and solutions.
  - Solves the problem of how concepts can figure into word meaning while keeping lexical items in the scope of semantics as a part of linguistic theory.
  - Gives us a way to approach some hard problems about word meaning and content.

Roots in the Adjectival Domain I

- Like to extend the packaging approach to other lexical categories.
- Consider adjectives here, focusing on gradable adjectives.
- Assume a degree analysis (Kennedy, 1997, 2007; Barker, 2002; Bartsch & Vennemann, 1973; Bierwisch, 1989; Cresswell, 1977; Heim, 1985; von Stechow, 1984).
- For example, the meaning of *tall* is given by a function to degrees on a scale:

(5) $[\![\text{tall}]\!](x) = d$ a degree of tallness

Roots in the Adjectival Domain II

- A scale is an ordered collection of degrees, with a *dimension* specifying what the degrees represent (e.g. Bartsch & Vennemann, 1973; Kennedy, 1997).
- Unpacking, the codomain is really a complex object $\langle D_\delta, <_\delta \rangle$, where:
  - D_δ is the set of degrees of dimension δ.
  - Ordered by $<_\delta$.
  - δ can specify e.g. tallness, speed, etc.
- So, a more explicit entry would be something like:

(6) a. $S_{\text{tall}} = \langle D_{\delta_{\text{tall}}}, <_{\delta_{\text{tall}}} \rangle$

b. $[\![\text{tall}]\!] : D_e \rightarrow S_{\text{tall}}$

Roots in the Adjectival Domain III

- Ultimately, want to claim that δ functions as the root for a gradable adjective.
- It is the dimension δ that does the work of providing the specific content of any particular adjective.
- But, need to explore the nature of scales a little more to illustrate this.
- Two general approaches to scales (cf. Solt & Gotzner, 2012).
    - Comparison based (e.g. Cresswell, 1977).
    - Abstract (e.g. Kennedy, 2007; von Stechow, 1984).
- For most semantic purposes, difference does not matter.
- But, slightly different takes on the root depending on which we opt for.

Roots in the Adjectival Domain IV

- Comparison approach (following Cresswell).
  - Begin with a relation *HIGHER* on individuals (suppressing domain presuppositions).
  - Define up equivalence classes on individuals $[a_e]_{HIGHER}$.
  - Then can set:

(7) a. $D_{\delta_{\text{tall}}} = \{[a_e]_{HIGHER}\}$

b. $a <_{\delta_{\text{tall}}} b$ iff b *HIGHER* a

  - The degree function for *tall* is given by:

(8) $[\![\text{tall}]\!](x) = d \quad \text{iff} \quad x \in d$

Roots in the Adjectival Domain V

- Abstract approach. One way of spelling out.
  - Start with a system of magnitudes, *HEIGHT*. Assume:

(9) a. $HEIGHT = \langle M, <_M, h \rangle$

b. M an appropriate system of magnitudes ordered by $<_M$.

c. h maps (appropriate domain of) individuals to values in M.

- Want $S_{\text{tall}} = \langle D_{\delta_{\text{tall}}}, <_{\delta_{\text{tall}}} \rangle$ to be a representation of these magnitudes.
- So, have:

(10) $h : HEIGHT \rightarrow \langle D_{\delta_{\text{tall}}}, <_{\delta_{\text{tall}}} \rangle$ a homomorphism

- Then can have:

(11) $[\![\text{tall}]\!](x) = d \quad \text{iff} \quad h(HEIGHT(x)) = d$

Roots in the Adjectival Domain VI

- Why *h* and where does *h* come from?
  - Some reason to think language uses particular sorts of scales, perhaps always dense linear orderings, with or without endpoints (Fox & Hackl, 2006; Kennedy & McNally, 2005). Countable??
  - If so, then ordering structure of any $\langle D_\delta, <_\delta \rangle$ is selected from a small menu (provided by FHL?).
  - Main role of *h* is just to select, and set up that it is the ordering linked to a particular system of magnitudes.
  - So, maybe existential quantification $\exists h$?
  - But also, only some magnitude systems will allow homomorphisms to these sorts of scales.

Roots in the Adjectival Domain VII

- Roots and packaging.
    - Either way, we have a root, either *HIGHER* or *HEIGHT*.
    - Either way, this is packaged into a scale, either by forming equivalence classes, or by constructing a representation of an abstract system of magnitudes.
    - This is then packaged into a degree-valued function.
    - That is the appropriate lexical entry for an adjective.
    - Again, range of packaging options seems to be linguistically determined. Assume to be part of FHL.
    - So, same structure of roots and packaging as with verbs.

Roots in the Adjectival Domain VIII

- The root remains atomic:
  - Semantics makes use of the topological properties of $\langle D_\delta, <_\delta \rangle$ determined by the ordering (e.g. Kennedy & McNally, 2005).
  - It also makes use of the structure of degree-valued functions, e.g. in comparative constructions.
  - Might be more structure for scales is required? See discussion of measure phrases in Sassoon (2010) or comparison classes in Solt & Gotzner (2012) (though for the latter, I am unsure how much is grammaticalized).
  - So, as before, we have some marking of structural features associated with the roots, but no features of how a concept is represented seems to be visible to grammar.

Roots in the Adjectival Domain IX

- Hence, treat the root as a pointer to an extra-linguistic concept, just as with verbs.
- Pointer will mark for structure, e.g. choosing HEIGHT or HIGHER, or other appropriate structure.
- But, need not see the internal structure of mental representation.
- Same as in the verbal domain.

Concepts and Degrees I

- A brief glance at some work on concepts for dimensions.
    - The roots for gradable adjectives need to provide dimensions like tallness, heaviness, speed, etc.
    - Treat as a pointer to a concept.
    - Ask how these might be represented.
    - Point out view in psychology that is suggestive of the abstract magnitude view (for some adjectives).

Concepts and Degrees II

- WEIGHT appears to be part of a commonsense physical theory (Carey, 1991, 2009).
  - Involves concepts like MATTER, WEIGHT, DENSITY.
  - Distinguishes material from immaterial things.
  - Weight a property of material thing, proportional to the quantity of matter.
- Suggests an abstract view of the degrees for *heavy*
  - At least, the appropriate concepts for a system of magnitudes seem to be available.
  - Not definitive. As with the case of stative verbal roots, might be that the lexical entry points to one of many available concepts.
  - But, makes the abstract approach natural.

Concepts and Degrees III

- Not clear this generalizes??
  - TALL has been argued (Keil & Carroll, 1980) to show a structure of *exemplars*.
  - Related to prototype structure.
  - But specific examplars for specific kinds of objects.
  - They suggest a single concept is abstracted from exemplars during development.
- Might suggest different sorts of scales for different sorts of adjectives (cf. Sassoon, 2010)?
- Or just different conceptual roots requiring different packaging mechanisms?

References I

Barker, C. (2002). The dynamics of vagueness. *Linguistics and Philosophy*, 25, 1–36.

Barsalou, L. W. (2003). Situated simulation in the human conceptual system. *Language and Cognitive Processes*, 18, 513–562.

Bartsch, R. & Vennemann, T. (1972). The grammar of relative adjectives and comparison. *Linguistische Berichte*, 20, 19–32.

——— (1973). *Semantic Structures: A Study in the Relation between Semantics and Syntax*. 2nd edn. Frankfurt: Athenäum.

Barwise, J. & Cooper, R. (1981). Generalized quantifiers and natural language. *Linguistics and Philosophy*, 4, 159–219.

References II

Bierwisch, M. (1989). The semantics of gradation. In M. Bierwisch & E. Lang (Eds.), *Dimensional Adjectives*, pp. 71–261. Berlin: Springer-Verlag.

Bierwisch, M. & Schreuder, R. (1992). From concepts to lexical items. *Cognition*, 42, 23–60.

Bloom, P. (2000). *How Children Learn the Meanings of Words*. Cambridge: MIT Press.

Bowerman, M. & Levinson, S. C. (Eds.) (2001). *Language Acquisition and Conceptual Development*. Cambridge: Cambridge University Press.

Carey, S. (1985). *Conceptual Change in Childhood*. Cambridge: MIT Press.

——— (1991). Knowledge acquisition: Enrichment or conceptual change. In S. Carey & R. Gelman (Eds.), *The Epigenesis of Mind: Essays on Biology and Cognition*, pp. 133–169. Hillsdale: Erlbaum.

References III

——— (2009). *The Origin of Concepts*. Oxford: Oxford University Press.

Chierchia, G. & McConnell-Ginet, S. (1990). *Meaning and Grammar*. Cambridge: MIT Press.

Clark, E. V. (1983). Meaning and concepts. In P. H. Mussen (Ed.), *Handbook of Child Psychology*, vol. III: Cognitive Development, 4th edn., pp. 787–840. New York: Wiley.

Cresswell, M. J. (1977). The semantics of degree. In B. Partee (Ed.), *Montague Grammar*, pp. 261–292. New York: Academic Press.

——— (1978). Semantics and logic. *Theoretical Linguistics*, 5, 19–30.

Davidson, D. (1967). Truth and meaning. *Synthese*, 17, 304–323. Reprinted in Davidson (1984).

References IV

——— (1984). *Inquiries into Truth and Interpretation*. Oxford: Oxford University Press.

Dowty, D. R. (1979). *Word Meaning and Montague Grammar*. Dordrecht: Reidel.

Dowty, D. R., Wall, R. E., & Peters, S. (1981). *Introduction to Montague Semantics*. Dordrecht: Reidel.

Field, H. (1972). Tarski's theory of truth. *Journal of Philosophy*, 69, 347–375.

Fox, D. & Hackl, M. (2006). The universal density of measurement. *Linguistics and Philosophy*, 29, 537–586.

Gamut, L. T. F. (1991). *Logic, Language, and Meaning*, vol. 2. Chicago: University of Chicago Press. 'Gamut' is a pseudonym for J. van Benthem, J. Groenendijk, D. de Jongh, M. Stokhof, and H. Verkuyl.

References V

Gentner, D. & Goldin-Meadow, S. (Eds.) (2003). *Language in Mind*. Cambridge: MIT Press.

Glanzberg, M. (2011). Meaning, concepts, and the lexicon. *Croatian Journal of Philosophy*, 31, 1–29.

——— (forthcoming). Explanation and partiality in semantic theory. In A. Burgess & B. Sherman (Eds.), *New Essays in Metasemantics*. Oxford: Oxford University Press.

Gopnik, A. & Meltzzoff, A. N. (1997). *Words, Thoughts, and Theories*. Cambridge: MIT Press.

Grimshaw, J. (2005). Semantic structure and semantic content in lexical representation. In *Words and Structure*, pp. 75–89. Stanford: CSLI Publications.

Hale, K. & Keyser, S. J. (1987). A view from the middle. Lexicon Project Working Papers 10, Center for Cognitive Science, MIT.

References VI

——— (1993). On argument structure and the lexical expression of syntactic relations. In K. Hale & S. J. Keyser (Eds.), *The View from Building 20*, pp. 53–109. Cambridge: MIT Press.

——— (2002). *Prolegomenon to a Theory of Argument Structure*. Cambridge: MIT Press.

Hampton, J. A. (1993). Prototype modles of concept representation. In I. van Mechelen, J. A. Hampton, R. S. Michalski, & P. Theuns (Eds.), *Categories and Concepts: Theoretical Views and Inductive Data Analysis*, pp. 67–95. New York: Academic Press.

Harley, H. (2007). The bipartite structure of verbs cross-linguistically. Ms.

Heim, I. (1985). Notes on comparatives and related matters. Unpublished ms, University of Texas, Austin, available at http://semanticsarchive.net/Archive/zc0ZjY0M/.

References VII

Heim, I. & Kratzer, A. (1998). *Semantics in Generative Grammar*. Oxford: Blackwell.

Higginbotham, J. (1988). Contexts, models, and meanings: A note on the data of semantics. In R. Kempson (Ed.), *Mental Representations: The Interface Between Language and Reality*, pp. 29–48. Cambridge: Cambridge University Press.

——— (1989a). Knowledge of reference. In A. George (Ed.), *Reflections on Chomsky*, pp. 153–174. Oxford: Basil Blackwell.

——— (1989b). Knowledge of reference. In A. George (Ed.), *Reflections on Chomsky*, pp. 153–174. Oxford: Blackwell.

Higginbotham, J. & May, R. (1981). Questions, quantifiers and crossing. *Linguistics Review*, 1, 41–79.

Jackendoff, R. (1990). *Semantic Structures*. Cambridge: MIT Press.

References VIII

Kamp, H. & Partee, B. (1995). Protype theory and compositionality. *Cognition*, 57, 129–191.

Keenan, E. L. & Stavi, J. (1986). A semantic characterization of natural language determiners. *Linguistics and Philosophy*, 9, 253–326. Versions of this paper were circulated in the early 1980s.

Keil, F. C. (1989). *Concepts, Kinds, and Cognitive Development*. Cambridge: MIT Press.

Keil, F. C. & Carroll, J. J. (1980). The child's acquisition of 'tall': Implications for an alternative view of semantic development. *Papers and Reports on Child Language Development*, 19, 21–28.

Kennedy, C. (1997). *Projecting the Adjective*. Ph.D. dissertation, University of California, Santa Cruz. Published by Garland, 1999.

References IX

——— (2007). Vagueness and grammar: The semantics of relative and absolute gradable adjectives. *Linguistics and Philosophy*, 30, 1–45.

Kennedy, C. & McNally, L. (2005). Scale structure, degree modification, and the semantics of gradable predicates. *Language*, 81, 345–381.

Klein, E. & Sag, I. A. (1985). Type-driven translation. *Linguistics and Philosophy*, 8, 163–201.

Larson, R. K. & Segal, G. (1995). *Knowledge of Meaning*. Cambridge: MIT Press.

Laurence, S. & Margolis, E. (1999). Concepts and cognitive science. In E. Margolis & S. Laurence (Eds.), *Concepts: Core Readings*, pp. 1–81. Cambridge: MIT Press.

Lepore, E. (1983). What model-theoretic semantics cannot do. *Synthese*, 54, 167–187.

References X

Levin, B. & Rappaport Hovav, M. (1995). *Unaccusativity: At the Syntax-Lexical Semantics Interface*. Cambridge: MIT Press.

——— (2005). *Argument Realization*. Cambridge: Cambridge University Press.

Montague, R. (1973). The proper treatment of quantification in ordinary English. In J. Hintikka, J. Moravcsik, & P. Suppes (Eds.), *Approaches to Natural Language*, pp. 221–242. Dordrecht: Reidel. Reprinted in Montague (1974).

——— (1974). *Formal Philosophy*. New Haven: Yale University Press. Edited by R. Thomason.

Murphy, G. L. (2002). *The Big Book of Concepts*. Cambridge: MIT Press.

Murphy, G. L. & Medin, D. L. (1985). The role of theories in conceptual coherence. *Psychological Review*, 92, 289–316.

References XI

Nguyen, S. P. & Gelman, S. A. (2002). Four and 6-year olds' biological concept of death: The case of plants. *British Journal of Developmental Psychology*, 20, 495–513.

Pietroski, P. M. (2003). The character of natural language semantics. In A. Barber (Ed.), *Epistemology of Language*, pp. 217–256. Oxford: Oxford University Press.

——— (2005). *Events and Semantic Architecture*. Oxford: Oxford University Press.

——— (2010). Concepts, meanings and truth: First nature, second nature and hard work. *Mind and Language*, 25, 247–278.

——— (2012). Semantic monadicity with conceptual polyadicity. In M. Werning, W. Hinzen, & E. Machery (Eds.), *The Oxford Handbook of Compositionality*, pp. 129–148. Oxford: Oxford University Press.

References XII

Pinker, S. (1989). *Learnability and Cognition*. Cambridge: MIT Press.

——— (1994). *The Language Instinct*. New York: HarperCollins.

Prinz, J. J. (2002). *Furnishing the Mind: Concepts and Their Perceptual Basis*. Cambridge: MIT Press.

Rappaport Hovav, M. & Levin, B. (1998). Building verb meanings. In M. Butt & W. Geuder (Eds.), *The Projection of Arguments*, pp. 97–134. Stanford: CSLI Publications.

Rips, L. J. (1989). Similarity, typicality, and categorization. In S. Vosniadou & A. Ortony (Eds.), *Similarity and Analogical Reasoning*, pp. 21–59. Cambridge: Cambridge University Press.

References XIII

Rosch, E. (1978). Principles of categorization. In E. Rosch & B. B. Lloyd (Eds.), *Cognition and Categorization*, pp. 27–48. Hillsdale: Erlbaum.

Rosch, E. & Mervis, C. B. (1975). Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology*, 7, 573–605.

Sassoon, G. W. (2010). Measure theory in linguistics. *Synthese*, 174, 151–180.

Slaughter, V., Jaakkola, R., & Carey, S. (1999). Constructing a coherent theory: Children's biological understanding of life and death. In M. Siegal & C. C. Peterson (Eds.), *Children's Understanding of Biology and Health*, pp. 71–96. Cambridge: Cambridge University Press.

Smith, E. E. & Medin, D. L. (1981). *Categories and Concepts*. Cambridge: Harvard University Press.

References XIV

Solt, S. & Gotzner, N. (2012). Experimenting with degrees. *Proceedings of Semantics and Linguisic Theory*, 22, 166–187.

von Stechow, A. (1984). Comparing semantic theories of comparison. *Journal of Semantics*, 3, 1–77.

Tarski, A. (1935). Der Wahrheitsbegriff in den formalizierten Sprachen. *Studia Philosophica*, 1, 261–405. References are to the translation by J. H. Woodger as "The concept of truth in formalized languages" Tarski (1983).

——— (1983). *Logic, Semantics, Metamathematics*. 2nd edn. Indianapolis: Hackett. Edited by J. Corcoran with translations by J. H. Woodger.

Vigliocco, G. & Vinson, D. P. (2007). Semantic representation. In M. G. Gaskell (Ed.), *Oxford Handbook of Psycholinguistics*, pp. 195–215. Oxford: Oxford University Press.

References XV

Westerståhl, D. (1989). Quantifiers in formal and natural languages. In D. Gabbay & F. Guenthner (Eds.), *Handbook of Philosophical Logic*, vol. IV, pp. 1–131. Dordrecht: Kluwer.

Wunderlich, D. (1997). Cause and the structure of verbs. *Linguistic Inquiry*, 28, 27–68.

Zimmermann, T. E. (1999). Meaning postulates and the model-theoretic approach to natural language semantics. *Linguistics and Philosophy*, 22, 529–561.