

---

# Lexical Semantics in Formal Semantics: History and Challenges

Barbara H. Partee

[partee@linguist.umass.edu](mailto:partee@linguist.umass.edu)

ESSLLI RefSemPlus Workshop, August 2016

# 1. Introduction

- The early important achievements of formal semantics that made it of interest to linguists were in compositionality, i.e. ‘the semantics of syntax’.
- The parts of the lexicon that received serious attention were the parts most directly relevant to compositionality.
- There was rapid progress on a wide range of function words or morphemes and the constructions in which they occur -- determiners, especially quantifiers, tense and aspect markers, plurality, negation, focus-sensitive words like *only* and *even*, comparative and superlative morphemes, and more.
- But for open-class content words, as is well-known, Montague assigned semantic types, but did not try to specify their meanings, considering their study an empirical matter outside of formal semantics. They were just treated as unanalyzed constants, as in formal logic, to be assigned values by interpretation in a model.

## Introduction, *cont'd.*

- Over the course of the history of formal semantics, there has been great progress on the study of some semantic properties of open-class words, including aspectual properties of verbs, context-sensitive properties of words like *local* and *enemy*, polarity items, functional and relational nouns, mass nouns, and more.
- These properties have sometimes been captured via meaning postulates – constraints on their interpretation in a model --, a practice used by Montague (from Carnap), sometimes via partial lexical decomposition, a practice introduced by David Dowty on the model of what was done in Generative Semantics.
- Lexical semantic research within formal semantics all concerns properties that are crucial in semantic composition; these lexical studies have been driven by the goals of compositional semantics.
- What has remained unanalyzed is the “remainder” of the content of open-class words; those are still treated by most formal semanticists as unanalyzed primitives.

## Introduction, *cont'd.*

- The issue of whether and how to try to add a fuller treatment of lexical content to formal semantics has a complex history which I will discuss, but was a side issue until very recently.
- It still draws less attention within theoretical circles than it does in the context of cognitive and computational perspectives.
- There are theoreticians who explicitly defend the idea of leaving the “remainder” of the meaning of lexical items unanalyzed linguistically.
- And the status of lexical meaning has figured in some major debates about semantic theory.
- Many of Chomsky’s critiques of formal semantics are based on a consideration of lexical items that have meaning but don’t seem to refer to anything (*chances* in *take chances*, etc). Chomsky seems skeptical of anything semantic beyond a syntactic level of “logical form”.

## Introduction, *cont'd.*

- And on the other side, Putnam's classic argument for the claim that "Meanings aren't in the head" was based on words that act like rigid designators, like *water* and other natural kind terms. Putnam argued for a strongly non-psychologistic basis for some kinds of lexical meaning – cases where there seems to be ONLY reference.
- Those who consider the lexicon outside the realm of formal (or other) semantics tend not to think of it as part of linguistics at all, but rather part of a more general cognitive domain, perhaps part of general knowledge, not so modularly encapsulated as 'truly linguistic' knowledge.
- But even if one agrees with Montague's initial supposition that lexical semantics is a separate empirical field of study, and whether one calls it "linguistic" or not, it would of course be a major advance if we could understand what human lexical understanding is like and how it interfaces with the construction of sentence meaning and discourse interpretation.

---

# Introduction, *cont'd.*

Plan of the rest of the talk:

- 2. Semantic competence issues. What are the goals of ‘doing semantics’? What counts as success? And is there a difference between describing a person’s language and describing their competence in that language? These issues loom largest for lexical semantics, I believe.
- 3. Some history. Lexical semantics before and in the beginnings of formal semantics. Approaches and disputes.
- 4. Some achievements of formal semantics in the realm of lexical semantics, empirical achievements that don’t require resolving foundational questions.
- 5. Meaning postulates – some of their appeal, and how they connect to some early theories and possibly to distributional semantics, and how they may avoid some theoretical pitfalls.
- (6. Four areas of theoretical progress)


## 2. Semantic competence issues

- Chomsky brought the notion of *competence* to the fore, and made it central to the goals of syntax, and by extension presumably to all of linguistics.
- But for semantics, especially lexical semantics, it's not obvious that the unconscious knowledge of the native speaker actually characterizes the language they speak, as it seems to do so neatly in syntax.
- In syntax, if two speakers differ in their internalized syntactic rules, then we say that they speak different idiolects. There is no such thing as 'not knowing' the syntactic rules of your language – what you know defines what your language *is*.
- The syntax of the 'language of a community' is perhaps a more problematic notion than the syntax of an idiolect, a kind of idealization, or else requiring an explicit account of variation.
- And in syntax, even in the lexicon there are no worries-in-principle – the lexicon is finite, and at worst can be memorized.

## Semantic competence issues, *cont'd.*

- In semantics, on the other hand, some of the biggest worries about ‘competence’ have concerned the lexicon
- Putnam 1975: “So theory of meaning came to rest on two unchallenged assumptions:
  - (i) That knowing the meaning of a term is just a matter of being in a certain psychological state.
  - (ii) That the meaning of a term (in the sense of ‘intension’) determines its extension (in the sense that sameness of intension determines sameness of extension.)

I shall argue that these two assumptions are not jointly satisfied by *any* notion, let alone any notion of meaning.”





## Semantic competence issues, *cont'd.*

- Putnam used his famous Twin Earth thought experiments to argue that nothing in the narrow psychological state of a language user determines that *water* is H<sub>2</sub>O rather than XYZ. Yet given the causal history of the term in English, *water* in fact picks out H<sub>2</sub>O in our language.
- And Putnam argued that even though he could not tell a beech from an elm, *beech* and *elm* in his language did not both just “mean” “some deciduous tree”, but rather picked out the natural kinds ‘beech’ and ‘elm’, by virtue of his being a part of a language community with a certain history.
- In both kinds of examples, the words pick out determinate kinds; they have a definite intension (rigid) and extension, but these are not determined by whatever ‘semantic representation’ may be in the speaker’s head, not by the speaker’s “unconscious knowledge”.
- There are *no comparable problems* in syntax!

## Semantic competence issues, *cont'd.*


- Putnam's arguments and related concerns about the lexicon led me to conclude in some earlier papers (1977-1982, around the time of the early interdisciplinary Cognitive Science interactions in the "Sloan years") that **we really don't know our language.**
- That didn't seem terrible, though, just interestingly different from the situation in syntax.
- I related it to the theory-dependence of lexical meanings, the role of 'experts' (for *beech-elm*, etc.), 'meaning holism', meaning change over time. It also relates to the much greater interpersonal differences we find in lexical semantics than in syntax or phonology.
- My conclusion was that our knowledge of meanings of many words we use is incomplete and underspecified. (In some but not all cases, the meanings themselves may be underspecified, partly 'open'.)
- For an example of such indeterminacy and change happening in (not only) English right now, consider *husband* and *wife*. (*Mary and I are wives* may now be ambiguous – cf ... *mothers* vs. ... *sisters*.)

## Semantic competence issues, *cont'd.*

- An alternative view: A historically major opposition point of view within linguistics (Jackendoff, Fodor): Meanings should be thought of as something like mental concepts, perhaps in some innate language of thought. This was a highly ‘internalist’ picture.
- Even though that view did not have the particular problems that model-theoretic semantics did, I worried that it completely neglected an important higher-order intention in language users, namely to speak the same language as each other. (Cf. David Lewis’s *Convention*). Our words don’t simply refer to ‘our ideas’.
- Go back to Putnam’s observation that he could not tell a beech from an elm, and hence had no “internal” “semantic features” or “concepts” that would differentiate them “in his head”. That did not mean that *beech* and *elm* were synonymous for him. He knew nothing “about” them, but having acquired the words as he did, they picked out different kinds. And he believed beeches and elms to be different kinds without being able to *describe* any difference.

## Semantic competence issues, *cont'd.*

- One promising resolution to the foundational tension between Chomskyan psychologism and formal semantics comes from recent arguments by Stalnaker and Burge about the need to revise what we take psychological realism to be.
- Putnam '75: “meanings ain’t in the head”.
- Stalnaker 1989: meanings **are** in the head, but like footprints are in the sand: the problem was in taking too narrow a view of “in the head”.
- “ ...a representational system is a system that is capable of being in a range of alternative internal states that tend to be causally dependent on the environment in a systematic way.”



## Semantic competence issues, *cont'd.*

- Similarly in the context of work on the philosophy of perception, on ‘social anti-individualism’, and on the origins of objectivity (2010), Tyler Burge has argued that Chomsky’s view of grammar “in the mind” underestimates the degree to which “the natures of many mental states constitutively depend on relations between a subject matter beyond the individual that has the mental states.”
- Stalnaker’s and Burge’s work suggests that we should combine work on individual semantic representations with much more research on how those representations come to have the external relational content that makes reference possible.



## Semantic competence issues, *cont'd.*

- Conclusion: The apparent tensions between formal semantics and the goals of a linguistic theory of the competence of the language user were only apparent.
- Perception involves a relation between perceiver and external objective reality. Perception is normally veridical but *can* be in error.
- Reference and truth conditions involve a relation between language user and external reality. Language is normally understood correctly but the language user *can* be mistaken.
- Language of course goes beyond objective representation; not all language is about the real world. No one suggests that truth-conditions exhaust all of meaning; but the puzzle of how to relate truth-conditional semantics to semantic competence is resolvable, and cognitive science is becoming the richer for it.

### 3. Lexical semantics in antecedents to formal semantics.

- Leibniz (1646-1716) articulated the goal of a “characteristica universalis” as part of his goal of designing a logically perfect language: one basic symbol for every basic concept, and everything built up compositionally from there. Of course never achieved, but that idea has great appeal and is frequently reborn.
- 20<sup>th</sup> century logicians: strictly compositional semantics, with primitive constant symbols at the base – they are assigned values in a model.
- As early linguists thinking about semantics, Fodor and Katz in the 1960s posited universal **semantic features** on analogy with phonology. Primitive ideas about compositionality – union of sets of features.
- Generative semanticists in the 60s and 70s added logical structure (based on 1<sup>st</sup> order logic) and a certain amount of semantic decomposition – e.g. CAUSE(BECOME(NOT(ALIVE))).

## Lexical semantics in antecedents, *cont'd.*

- Anna Wierzbicka is probably the most persistent proponent of a set of **universal semantic primes** from which all lexical meanings can be composed, and she has done a lot of interesting typological work on that basis. But in general that approach is not considered likely to succeed as a general theory of lexical meaning.
- But see Avery Andrews's interesting work in progress, Andrews, Avery. 2015. Reconciling NSM and Formal Semantics (v4). Ms.  
[http://www.semanticsarchive.net/Archive/2YxMmlwN/NSM\\_Formal\\_Semantics\\_v4.pdf](http://www.semanticsarchive.net/Archive/2YxMmlwN/NSM_Formal_Semantics_v4.pdf).





## Lexical semantics in antecedents, *cont'd.*

- Carnap introduced meaning postulates as non-logical axioms that could guarantee the validity of various intuitively valid arguments that didn't follow by logic alone.




## Lexical semantics in antecedents, *cont'd.*

- Montague included some meaning postulates in his work, constraining which logically possible interpreted models could be models for English. I will say more about that approach in Section 5.



## Lexical semantics in antecedents, *cont'd.*

- Philosophers' reactions to linguists' "semantic representations" as a way of doing semantics – David Lewis (1970):
- “But we can know the Markerese translation of an English sentence without knowing the first thing about the meaning of the English sentence: namely, the conditions under which it would be true. Semantics with no treatment of truth conditions is not semantics.”
- “Translation into Markerese is at best a substitute for real semantics, relying either on our tacit competence (at some future date) as speakers of Markerese or on our ability to do real semantics at least for the one language Markerese.”



## Lexical semantics in antecedents, *cont'd.*

- But linguists *did* presuppose tacit competence in Markerese; they took it – or some kind of representation language -- to be universal and innate, and many still do (e.g. Jackendoff; also Jerry Fodor).
- To philosophers and logicians doing formal semantics, the language of Markerese looked empty, since it was uninterpreted.
- To linguists in 1970, concern with *truth* looked puzzling. Linguists were trying to figure out mental representations that would underlie linguistic competence.
- When the linguistic relevance of truth conditions finally penetrated, the nature of linguistic semantics changed, in terms of tools used, questions asked, and criteria of adequacy for semantic analyses.
- Looking forward: possible reconciliation if evolution can provide some of the external referential links needed to give ‘innate semantic features’ the kind of external grounding Lewis demanded, and if learning mechanisms can be found to account for the external content of non-innate semantic features or morphemes.

## 4. Empirical achievements of formal semantics in lexical semantics

- Formal semantics has actually made a great many contributions to lexical semantics, even without any general theory of lexical semantics.
- As noted in the introduction, the emphasis has been on points of interface between lexical content and compositionality. And what has been studied has been certain aspects of lexical meaning, not lexical meanings as a whole.
- As a prime early example, take the study of verbal aspect, starting with early work by Dowty, Bennett and Partee, Bach, Parsons.
- The driving challenge was to account for the compositional interpretation of sentences including various tenses and temporal adverbials, including entailments that had been catalogued by Zeno Vendler, such as *John is walking* → *John has walked* but not *John is building a house* → *John has built a house*.

## Empirical achievements, *cont'd.*

- The different entailment patterns between *John is walking* and *John is building a house* evidently reflect the different *Aksionsarten* of the verbs *walk* and *build*: one is a process verb, the other an accomplishment verb.
- Early work analyzed these aspectual *properties* of verbs of different *Aksionsarten* – states, processes, achievements, accomplishments, and was able to offer accounts of much of their compositional behavior with tenses, auxiliary verbs, temporal modifiers.
- At first all in terms of times – moments and intervals – without taking events as primitives in the ontology. But already much progress on semantics of perfect, progressive, *for an hour* vs. *in an hour*, etc.
- For instance, suppose *build a house* is true of an interval, and *be building a house* true at any moment in the interior of such an interval, and *have built a house* true at any moment following such an interval. That's enough to explain the non-entailment we noted between *John is building a house* and *John has built a house*.

## Empirical achievements, *cont'd.*

- Great progress in ensuing decades. Incorporating events (eventualities) into models as well as times (Parsons, Bach). Analyzing telicity, (im)perfectivity, and distinguishing them (Filip). Incorporating Reichenbachian distinctions among event time, reference time, and speech time (W. Klein).
- Shortly after Link and Sharvy revolutionized the study of the mass-count distinction in nouns and the semantics of plurals, their algebraic perspective was extended to the event domain by Bach and by Link. This 'subatomic' (Parsons) analysis of entities and events opened up more fine-grained semantic analysis of nouns and verbs – while still leaving unanalyzed the 'lexical remainder' distinguishing lexical items whose nominal or aspectual profiles were the same. No analysis of *dog vs cat* or *walk vs run*.
- But those advances in turn led to new insights into the interaction of aspect and quantification involving verbs, their arguments with various thematic roles, and temporal modifiers (Dowty, Krifka, Filip).

## Empirical achievements, *cont'd.*

- In recent decades, work on fine-grained event structure in the semantics of tense, aspect, and temporal modification has been enriched by typological work and by exciting work on language acquisition and language processing, reinforcing the ‘psychological reality’ of the posited semantic properties.
- In those domains and others, progress in lexical semantics and progress in compositional semantics have gone hand in hand.
- And none of it has been impeded by the lack of an account of lexical semantics ‘in general’.
- All of it, I believe, has been pretty much consistent with either of two approaches I mentioned in the introduction:
  - studying “semantic properties” of lexical units (captured for instance by meaning postulates), or
  - doing “partial decomposition”, identifying ‘functional heads’ or ‘operators’ with specified properties, applied in a structured way to the (unanalyzed) meaning of a given root lexical item.



---

## Empirical achievements, *cont'd.*

- I have focused on formal semantic achievements concerning verbs and nouns, and could also mention much work on adjectives, because this work concerns the open-class vocabulary, the part of the lexicon where formal semantics has the least to say, and no general theory.
- I've done this partly in reaction to my own field: I have often worried about the absence of a theory of lexical meaning, and have worried about the tendency of formal semanticists to endorse Montague's practice of leaving lexical semantics outside of 'semantics proper', possibly outside of 'linguistics proper' (and in that respect to partially side with Chomsky's skepticism towards 'real semantics'.)

## Empirical achievements, *cont'd.*

- But I shouldn't neglect to at least mention some of the achievements on formal semantics's 'home turf', such as the semantics of quantifiers (Montague, Barwise and Cooper, Keenan, ...), and one of my favorites, the semantics of cross-categorial conjunction (Dowty, Gazdar, Partee and Rooth).
- The semantics of cross-categorial conjunction – *and* and *or* connecting sentences, DPs, VPs, AdjPs, and more – has turned out to be systematic and universal (universal allowing for a finite number of marked exceptions in any language).
- In this case the semantic analysis of the given lexical items is complete analysis: these are closed-class items whose content is entirely “logical”.
- Montague included analysis of *and* connecting S's and VPs (not of NPs only because he wasn't ready to deal with plurals). Dowty et al generalized what Montague had done to cover all 'Boolean conjoinable categories'.

## Empirical achievements, *cont'd.*

- Upshot: Conjunction seems to be a universal and distinct “module” of the grammar. Both in syntax and in semantics its properties seem to be special and universal.
- Of course even with conjunction the story is not complete. “Or” is more puzzling than “and” (Simons) and work on it continues. Interaction with negation and other scope-bearing elements is non-trivial. The relation between ‘Boolean’ *and* and ‘group-forming’ *and* is also non-trivial (Winter).
- But the cross-categorial semantics of Boolean *and* is well-enough understood and robust that it can be employed as one test for evaluating analyses of various constructions, as for instance to argue in favor of  $\langle e,t \rangle$  type for DPs in predicative positions, where they can be conjoined with APs. (*John is intelligent and an authority on unicorns.*)
- Conclusion: Formal semantics has made contributions to aspects of lexical meaning, even without a theory of lexical meaning.

## 5. About meaning postulates

- The use of meaning postulates in Montague's work seems somewhat *ad hoc*, and rather few semanticists have advocated such an approach in general (but see Borschev and Partee 1998.)
- Ede Zimmermann in a 1999 paper demonstrated the danger of trying to use them to get around certain kinds of compositionality problems, showing that they must be used with caution.
- But I believe that meaning postulates can be a very useful resource for lexical semantics for at least two reasons.
- **First**, from within the perspective of formal semantics, with the goal of doing something about lexical semantics. If you don't believe in, or are agnostic about, the existence of some set of universal semantic primitives, then the possibility of actually giving complete and non-circular definitions for all words is impossible. Similarly, insofar as the interpretation of some words is presumed to be at least in part a matter of reference (cf. *beech* and *elm*), we may lack a metalanguage in which to "give" that reference.

## About meaning postulates, *cont'd.*

- But meaning postulates are a way to constrain, i.e. partially specify, lexical meanings without claiming to provide definitions.
- Montague didn't use the word 'meaning postulates' for the ones he provided – he simply called them 'constraints on possible models' that were to count as models of English. Without them, any lexical item whose interpretation was not spelled out explicitly (for him, those were *be, and, or, not, will, have –en, necessarily*, and perhaps a few more) could in principle be given any model-theoretic interpretation of the appropriate semantic type.
- Some of his meaning postulates specified logical properties of some lexical items. E.g., in the general case, all argument-taking expressions took intensional arguments -- for example, all transitive verbs were given the semantic type appropriate for *seek*, taking the intension of a generalized quantifier as object. Then he used meaning postulates to state that some of the transitive verbs have meanings that amount to first-order extensional relations.

## About meaning postulates, *cont'd.*

- Other meaning postulates were of a sort very relevant to the general issue of lexical semantics. These were meaning postulates that established word-to-word or word-to-phrase connections, guaranteeing the non-independence of various word meanings.
- An example from PTQ is a meaning postulate guaranteeing the equivalence of *seek* and *try to find*, which he used in order to show the equivalence of his treatment of *seek* as a basic transitive verb with a Quinean decompositional treatment of *seek* into *try to find*.
- One can also use meaning postulates to specify lexical entailments of the “IS A” sort among subordinate-superordinate predicates, or converse relations among words like *employer*, *employee*, or to specify that some relational words like *cousin* are symmetrical.
- Assuming that the metalanguage for expressing meaning postulates is rich enough to distinguish presuppositions from entailments, they can also be used to specify the sortal requirements of argument-taking words, given that those are presumably presuppositional.

## About meaning postulates, *cont'd.*

- In general, almost any clauses that might be part of a lexical entry in an ordinary monolingual dictionary could be rephrased as meaning postulates, constraining the interpretation of the head word to be compatible with the interpretation of the given clauses.
- Thus meaning postulates can be viewed as something like “modest partial definitions”, modest in two ways: (i) They don’t pretend to be exhaustive or complete definitions, and (ii) they generally connect words to words (or sometimes to logic), not to external reality.
- And one would want a metalanguage for meaning postulates rich enough to include indexical and demonstrative elements, so as to be able to say, for instance, that *this afternoon* is part of *today*, etc.

## About meaning postulates, *cont'd.*

- The **second** reason that meaning postulates may be useful is especially relevant in the context of this workshop. Distributional semantics, from the little I know about it, does not try to say anything about reference, but treats lexical semantics in terms of the distribution of words in texts, going back to Zellig Harris's observation that similarity of cooccurrence generally corresponds to similarity of meaning (as well as to similarity of syntactic category).
- Meaning postulates, at least prototypical ones, also relate words to words. They might thus be able to provide some sort of bridge between formal semantics and distributional semantics. I don't think it would be a very direct bridge, given that distributional vectors involve numbers that specify some kind of similarity metrics, and familiar sorts of meaning postulates do not say anything about 'similarity'.



## About meaning postulates, *cont'd.*

- But both meaning postulates and distributional semantics relate expressions to expressions, constraining reference without actually specifying reference.
- The methods of both can, I think, be easily adapted to describe specialized sublanguages, e.g. different dialects, or sublanguages characterized by technical jargon or the like. Meaning postulates are at heart just axioms, and it is easy to imagine modeling the difference between one dialect and another, or a general vs. a technical context, by differences in some choices of axioms. In the case of distributional semantics, I suppose that such versatility is built in, simply via the choice of data set to use as input.
- But since I know so little about distributional semantics, and since the range of kinds of meaning postulates imaginable is so open-ended, these remarks can only be suggestive and speculative.
- Vladimir Borschev and I have appealed to meaning postulates in some of work endeavoring to build bridges between the rich Russian tradition of lexical semantics (esp. Apresjan) and western formal semantics, showing how to recast Apresjan's 'lexical portraits' in terms of meaning postulates.

## 6. Progress in theorizing about the lexicon and semantic competence

- If time remains, I want to close by mentioning four areas where formal semanticists and philosophers have made valuable contributions to theoretically important problem areas in formal semantics. These are just thumbnail mentions, to possibly stimulate discussion, and far from exhaustive.
- On the problem of vagueness, one that is crucial to lexical semantics but hardly arises for compositional semantics: important work early on by Kamp, Pinkal, and others. Supervaluation theory is one approach, various approaches to underspecificity. Work on prototype theory is directly relevant here.
- Context-dependence – formal semanticists have helped to uncover how widespread context-dependence of lexical meaning is – not just in indexical and demonstrative words, but in words like *local*, *neighbor*, *right* and *left*. Several major approaches, with and without ‘invisible pronoun’-like elements and parameters of interpretation.

## 6. Progress in theorizing, *cont'd.*

- Prototype theory – formal semanticists have contributed to progress by helping to clarify when prototype-like meanings are part of semantics proper (*red, chair*) and when words may have relatively sharp meanings but prototype-like associated stereotypes (*grandmother*), and by suggesting tools for extending compositional semantic tools to deal with meanings that correspond to prototypes.
- The psychological reality puzzles – can meanings be in the head, etc.? As noted in Section 2, major advances on that front have come from philosophers in recent decades, especially Stalnaker and Burge, in work that is more philosophy of mind than philosophy of language proper. The crucial insight is that psychological states should not be thought of in the excessively ‘internalist’ way; psychological states should be conceived relationally,

# Selected references

More material and fuller references can be found in several papers, versions of which are downloadable from my site, <http://people.umass.edu/partee/> . Fuller references for things I referred to in this talk are available on request.

- Borschev, Vladimir, and Partee, Barbara H. 1998. Formal and lexical semantics and the genitive in negated existential sentences in Russian. In *FASL6: The Connecticut Meeting 1997*. <http://people.umass.edu/partee/docs/russian.pdf>.
- Burge, Tyler. 2010. *Origins of Objectivity*. Oxford.
- Partee, Barbara H. 2011. Formal semantics: Origins, issues, early impact. In *Formal Semantics and Pragmatics. Discourse, Context, and Models.*, eds. Partee, Glanzberg and Skilters, 1-52. <http://thebalticyearbook.org/journals/baltic/article/view/1580/1228> .
- Stalnaker, Robert. 1989. On what's in the head. *Philosophical Perspectives* 3: Philosophy of Mind and Action Theory:287-316.