

PARADOXES AND SEMANTIC REPRESENTATION

Richmond H. Thomason
Linguistics Department
University of Pittsburgh
Pittsburgh, PA 15260

ABSTRACT

Many researchers in Computer Science, Linguistics, Logic, and Philosophy have been discovering in various ways that analogues of the Liar paradox pose deep foundational problems for intensional semantics. In this survey paper I state my own view of what the problems are, and try to provide a broad perspective on the issues, with many references to the literature. Although the paper concentrates on developments in Philosophical Logic, I hope that the paper will help all researchers concerned with these problems to locate their concerns in a broad, interdisciplinary framework.

1. The Liar Paradox

The problematic traits of the Liar Paradox are hardy, and likely to survive transplanting. We have become interested in what can be learned from intensional variants of the Paradoxes, variants that apply to propositional attitudes in a computational setting. But it is worth beginning with problems posed by logical paradoxes in general, and in particular by the Liar.¹

A paradox is a plausible argument leading to an implausible conclusion. Paradoxes have always intrigued logicians, and many of the most interesting results of modern symbolic logic have been obtained by showing how certain paradoxes, or near-paradoxes, can be precisely formalized.

When we reach an implausible conclusion, we feel compelled to revise our opinions; so paradoxes call out for explanatory solutions: a satisfactory account of what went wrong in the reasoning. In case the paradox can be formalized (say, in the form of an inconsistent theory), there will be a related technical project of making reasonably small changes in the formalization that issue in a consistent theory (or at least, in a theory not known to be inconsistent).

The Liar Paradox seems to have been designed to humble logicians. Tarski's theorem,² after all, is a theorem, and imposes constraints on attempts to solve the Liar, just as Gödel's theorem imposes constraints on programs aimed at providing foundations for mathematics. This forces a technically satisfactory solution to the Liar to adopt at least one of the following three strategies.

- Strategy 1:** Hold certain seemingly intelligible notions to be inexpressible. (A special case of this would be giving up classical logic—holding classical negation, for instance, to be inexpressible.)
- Strategy 2:** Give up plausible schematic principles on truthlike predicates, such as Convention T.
- Strategy 3:** Impose limits on the extent to which language can be used to talk about its own syntax.

¹In keeping with modern treatments of the Liar, I am treating it as a phenomenon of *direct discourse*; the “what” in ‘What I am saying is false’ is a sentence. We are leading up to variations on the Liar in which the “what” is construed as a proposition, a belief or the like.

²See [Tarski 1956].

The classification is meant to divide approaches to the semantic paradoxes along lines corresponding to the pressures that will apply in attempting to make a case that the approach represents an explanatory solution. A language-limiting strategy, for instance, will have to argue that some seemingly sensible notions in fact are illusory; or at least, can't be expressed. A principle-limiting strategy must argue that the generality of the relevant principles is in fact less than it seems to be. And a self-reference-limiting strategy will have to argue that, for some reason, things that really are language-like can't be described as such—at least, not in full detail.

The language-limiting strategy includes solutions that introduce language-internal hierarchies, like the ramified theory of types,³ as well as those, like Tarski's metalinguistic solution, that make the hierarchies language-external. Putting the hierarchy into linguistic context of utterance, as Burge⁴ suggests, is an interesting compromise between the internal and external treatments. Still, like the other hierarchical solutions, it has the consequence of making many things unexpressible; in particular, there is no type-free way to say that a sentence is true. Thus, the *explanation* that is required in order to turn a technical solution of this sort into an explanatory solution of the paradox will be an argument that certain things that seem to be sayable in a type-free way in fact are not.

These hierarchies can all be extended naturally into the transfinite ordinals, and at successor ordinals seem to involve a process in which a process of reflection on a lower level somehow brings to light new truths.

With a certain amount of reluctance, I decided to include cases in which classical logic is given up under the first heading as well. This agrees well with solutions, like Kripke's,⁵ that appeal to truth-value gaps and "choice negation." In explaining such theories, it is natural to fall into a turn of speech in which the sentences that are not true are divided into (i) those whose negations are true, and (ii) all the rest. This seems to create a usage of 'not true' that is not represented in the theory itself. So, if you want to show an approach of this sort to be explanatory, you'll have to argue that the sense of negation that seems to correspond to this usage is illusory.⁶ These considerations would also apply plausibly to the type-free lambda calculus as an approach to the Liar. Here, in view of the Curry paradox,⁷ a conditional satisfying *modus ponens* must be ruled out of linguistic bounds, as well as a classical negation.

³[Whitehead & Russell 1910]

⁴See [Burge 1979].

⁵See [Kripke 1975].

⁶Kripke's paper makes it clear that he is well aware of this problem.

⁷See [Curry 1942].

The logical solution that seems to fit the language-limiting pattern least well is a paraconsistent approach, based on a version of relevance logic.⁸ The tactic that enables theories to be nontrivial theories to be inconsistent seems to depend less on restricting the background language than on giving up the principle of (material) *modus ponens*:

$$[\phi \wedge (\sim\phi \vee \psi)] \rightarrow \psi$$

For this reason, it might be best to place paraconsistent approaches with principle-limiting strategies. But since—for reasons that are rooted fairly deep in the philosophy of language—it is always difficult to distinguish between giving up *principles of logic* and limiting ones language, I prefer to classify it under the first category, with the other logical solutions, while remembering that it is exceptional in some ways.

Gupta's approach⁹ is a good example of the principles-limiting strategy. The application of jump operators and fixpoints to the technical analysis of truth derives from Kripke, but is placed in a setting that preserves classical logic. Something else has to give—and what gives is Convention T.

The self-reference-limiting strategy has at least the following technical interest: how much can be preserved of the apparatus of syntax if enough syntax is given up to block proofs of Tarski's theorem?¹⁰ Gupta's results¹¹ show that a surprising amount can be done with a restricted syntax of this sort.

On the other hand, the explanatory potential of the self-reference-limiting strategy seems to be unpromising. An explanation would have to make it plausible that, although our theory uses genuine names of sentences, it will become inconsistent if it is made explicit that this is what they name. But it is very hard to see how merely adding more linguistic detail about things we are already describing in some dim fashion could create contradictions. Moreover, as many people have noticed and Kripke shows vividly in [Kripke 1975], it isn't necessary to use an indirect construction like Gödel's diagonal method to construct self-reference of the sort demanded by the Liar; the ability to baptise things with a proper name also serves the purpose.

Commentators on [Kripke 1975] have neglected what seems to me to be one of the most important contributions to the recent literature on the Liar. This is the point,

⁸See [Priest 1979], [Visser 1984], and [Woodruff 1984]

⁹See [Gupta 1982] and [Gupta 1984].

¹⁰We know from Gödel's work that if a syntactic theory is strong enough to enable every primitive recursive syntactic function to be characterized, then it will enable us to prove Tarski-like theorems, *via* Gödel's diagonalization technique.

¹¹See [Gupta 1982].

which Kripke establishes by artful deploying of examples, that contingent, extralinguistic considerations can play a rôle in determining whether a sentence is used paradoxically. Any solution that does not show how this can occur is inadequate.

2. The Liar in an Intensional Setting

The Liar Paradox shows that truth is somehow *ineffable*; the relation between sentences and their truth values is problematic, and theories that enable this relation to be expressed are subject to various maladies.¹²

Intensionality has to do with semantic representations of sentences that are more full-blooded than mere truth-values, and which therefore are able to support interpretations of modal operators, or perhaps of propositional attitudes like 'know', 'believe', and 'be likely'. The general way of dealing with such constructions is to interpolate certain values between sentences and their truth-values; I will call these values *propositions*, without assuming anything about what propositions are.

Propositions are the middle-men between sentences and truth; sentences express propositions, and propositions have truth values. As long as there are enough propositions—so that, for instance, if one sentence is contingent and another necessary we don't have to associate the same proposition with the two—we have the materials to provide formally correct interpretations of intensional constructions.

Explanatory interpretations would require propositions to have more structure. In terms of this structure, for instance, if we could define operations of conjunction, necessitation, and belief on propositions we could then let the conjunction of ϕ with ψ be the conjunction of the proposition expressed by ϕ and the proposition expressed by ψ ; the necessitation of ψ be the necessitation of the proposition associated with ψ ; and the "belief" of ϕ be the "belief" of the proposition associated with ϕ .¹³ The explanation, then, consists in establishing a semantic operation, like necessitation on propositions, and then using it to interpret a syntactic operation.

However, Tarski's theorem applies to languages with intensional operators; and truth will still be ineffable in such languages. And if truth is a composition of two relations, one of them must be ineffable. This means that either (i) the relation between sentences

¹² Anyone who is worried about my not taking into account the fact that this relation is sensitive to context of utterance should suppose that what I say is relative to a fixed context.

¹³ The possible worlds approach to propositions accomplishes this by arranging things so that the propositions are members of a closure algebra. Boolean connectives map into boolean operations of the algebra; necessitation maps into the interior operation.

and the propositions they express or, (ii) the relation between propositions and their truth values, must be problematic.

Both alternatives are uncomfortable. On the first, we make the semantic interpretation of sentences mysterious; on the second, we lose the ability to deploy certain kinds of semantic explanations. The very simple explanations that truth-functions offer for boolean connectives like conjunction, for instance, would precipitate a hierarchy of metalanguages, because the notion of the truth of a proposition for a language L would not be expressible in L .

The difficulty is exacerbated by the Montague-Löb results,¹⁴ which stand to the Knower Paradox as Tarski's theorem stands to the Liar Paradox. Tarski's theorem shows that a truth connective T , satisfying the scheme

$$T\phi \leftrightarrow \phi,$$

will produce inconsistency if it is construed as a syntactic predicate. Montague's results show that connectives satisfying very weak modal conditions will do the same.

Such results have the effect of enlarging the list of natural language constructions that will be mysterious or unexplainable if we follow alternative (ii) and make the relation between sentences and the propositions they express unproblematic. Not only words like 'necessarily' and 'possibly' satisfy these modal conditions, but also the modal auxiliaries, and words like 'provable'. Also, it is convenient to place modal conditions on epistemic attitudes like 'know'; either as an idealization, or as a formulation of "implicit knowledge."¹⁵ Under these idealizations, semantic explanations of epistemic attitudes also will be threatened.¹⁶

Hilary Putnam argues that there is a tension between two common assumptions about meaning.¹⁷

- (I) To know the meaning of a term is to be in a certain psychological state.
- (II) The meaning of a term determines its extension (and as a special case, the meaning of a sentence) determines its truth value.

The difficulty that I have been trying to develop can be put as a somewhat similar tension between assumptions.

¹⁴See [Löb 1955] and [Montague 1974].

¹⁵The crucial part of the modal idealization consists of conditions guaranteeing that $K\phi$ will hold for each logically valid first-order sentence ϕ .

¹⁶[Thomason 1980] extends the Montague results to idealized belief-like connectives, which do not satisfy the scheme $B\phi \rightarrow \phi$.

¹⁷I paraphrase from pp. 135-6 of [Putnam 1975].

- (I') The semantic representations of sentences are *graspable*, in a sense that implies that they can be calculated.
- (II') The semantic representations of sentences are *alethic*, in a sense that implies that they yield relatively simple explanations of boolean connectives, modal operators, and the like.

This pair of conflicting assumptions is somewhat similar to Putnam's. But it seems to me that the second opposition raises deeper problems for semantic theory, because the considerations that fuel the conflict are grounded in logical limitations related to the semantic paradoxes, and because graspability is a clearer and more tractable concept than psychological representability.

Since a semantic theorist who stresses the desirability of (I') is likely to favor a view of semantic interpretation as translation or compiling, and one who stresses the desirability of (II') is likely to favor explanations having to do with truth and reference, the opposition is related to representationalist-realist debates in the foundations of cognitive psychology.¹⁸

But again, an opposition related to the paradoxes puts the matter in a different light. It suggests a less controversialist approach to the issues, since for technical reasons each of the two alternatives will be surrendering something desirable and a synthesis combining the merits of the two approaches will not be readily forthcoming.

Finally, the opposition between (I') and (II') is closely related to *Benacerraf's problem* in the philosophy of mathematics. Benacerraf's opposition is between finding a philosophical account of mathematics that is both epistemologically satisfactory, and capable of giving a plausible ontology to classical mathematics. Stressing the need to explain how we can have knowledge in mathematics leads to programs that cross the border from philosophy of mathematics to mathematics itself, since the need to provide such explanations is closely related to the revisionist programs in mathematics associated with constructivism. The need to be true to classical analysis makes it very hard to put mathematics in a reasonable relation to a knowing subject. It seems to me that Benacerraf's problem is best regarded not as distinct from my opposition between graspability and explicability in semantics, but as the special case of the difficulty where mathematical language is at stake.

This approach to foundational difficulties in semantics suggests that in coming to grips with the problems, we must look to work on technical solutions to the Liar Paradox. Since most of the recent outpouring of work on this topic has been motivated by a pure interest in the challenge of the semantic paradoxes, this seems to be another case of mathematics coming into being as it is needed, by a kind of preestablished harmony. Occasionally when

¹⁸See [Fodor 1975] and [Dennett 1978]; also, contrast [Fodor 1975] with [Stalnaker 1984] and [Barwise & Perry 1983].

I have been present at talks on the paradoxes, someone in the audience has asked what the wider relevance of the technical work is. Generally, the answer that I have heard is that the work is interesting, regardless of questions of relevance. This is true, but I believe that a much stronger case can be made for the usefulness of this work.

3. Ways out

Placing the Liar in an intensional setting does not seem to open many solution strategies that cannot be classified as modifications of approaches to the Liar. Each of the three general strategies canvassed in Section 1 could be applied to the Knower Paradox, and—depending on the extent to which the solution is explanatory—could be made into a foundation of a theory of intensional semantic representation, in natural language semantics or computational semantics.

There is still much work to be done in developing the details of the necessary theories. This is becoming an active area of research, judging from the projects that I'm aware of,¹⁹ after the dust has settled a bit, it should be possible to make comparative judgments with more confidence; nevertheless, it looks as if the pattern of technical solutions that will emerge will be similar to ones that have been developed for the Liar.²⁰

Against this background, the new approaches to the problem that an intensional setting affords stand out as particularly noteworthy. I see three such possibilities.

¹⁹See [Perlis 1985], [Asher & Kamp 1986], [des Rivières & Levesque 1986], [Kremer 1986], and [Morgenstern 1986]. Since time was short for preparing this paper, I have not said much about research on type-free programming languages, and work in natural language semantics inspired by this research. However, this research is very relevant, especially as a source of insights into the positive constraints that need to be placed on a successful theory. That is, we can see from this work what expressive limitations would be too severe to enable certain programming strategies to be implemented naturally, or to enable sense to be made of natural language constructions that seem perfectly sensible from an intuitive standpoint. See [Chierchia 1982], [Stoy 1977], and [Reynolds 1983].

²⁰In some cases, of course, the relation is explicit and intended; Asher and Kamp's approach, for instance, is a modification of Gupta's; Kremer and Morgenstern base their work on Kripke's construction.

- Strategy 1': Impose limits on the extent to which language can be used to talk about its propositions, while at the same time allowing it the full expressive power of quotation.
- Strategy 2': Do not require every sentence to express a proposition, and explicitly limit Convention T to sentences that express propositions.
- Strategy 4: Claim that propositional attitudes are not alethic, do not satisfy modal conditions. Thus, though Liar-like puzzles may arise with truth and necessity, they could not arise with attitudes, like knowledge and belief.

Strategy 1' depends on the fact that propositions—beliefs, claims, and the like—are somehow less tangible than sentences. Thus, the claim that their properties are problematic may be more defensible than the claim that syntactic properties are mysterious. I do not know of any official attempt to develop this strategy, but there is an intriguing relation to Benacerraf's solution of Lucas' puzzle about Gödel's incompleteness theorem. Essentially, Benacerraf's approach²¹ is to say that the ideal mathematician can know *that* he is a Turing machine, but cannot know *which* Turing machine he is. The analogous strategy in resolving the conflict between the graspability and the alethic character of propositions would be to claim that propositions are in fact sentences of the "language of thought," but that we cannot—on pain of inconsistency—know which sentences they are. This strategy has not, as far as I know, been developed, or even proposed; but it strikes me as an intriguing idea, worth following up, especially since much of the necessary mathematical work has been done by logicians exploring the consistency of Benacerraf's proposal.²²

Strategy 2' is suggested in [Thomason 1982]. However, there is a price: because of variations on the usual paradoxical constructions, of the sort discussed in [Parsons 1974], it is possible in this variation to prove some things that can also be proven not to express a proposition. The resulting theories are problematic in some ways, even though they need not be formally inconsistent.

Strategy 4 is defended in [Cresswell 1985]. The approach is interesting, but I am not convinced that it is explanatory. The best way of motivating this approach would be by means of linguistic examples showing that the modalities and the attitudes take different objects. Unfortunately, the evidence is mixed: 'What he believes is true' is perfectly acceptable. However, as Cresswell points out in [Cresswell 1985], p. 40, 'Jeremy believed

²¹See [Benacerraf 1967].

²²See [Reinhardt 1983] and [Flagg 1984].

that the sentence Miriam uttered and what Mary hinted were equally true' does not seem to be totally unintelligible, though it conflates a sentential and a propositional object.

What makes this strategy particularly difficult to defend, it seems to me, is that its proponent has not only to explain why the attitudes do not satisfy alethic conditions, but why it is a *logical* matter that they do not satisfy such conditions. It is easy enough for me to see that my own beliefs aren't closed under logical consequence. It is harder to see why if a rational agent's beliefs *were* closed under logical consequence, and other plausible conditions were met, a contradiction would ensue.

There remains one interesting strategy that I omitted from the list, because it is not new, and also is too complex to classify in a few words. This is Ramified Type Theory. This theory was developed with the intensional paradoxes in mind, and it provides a way of escaping not only the Knower Paradox, but other puzzles as well.²³ Russell and Whitehead gave the approach up in favor of the Simple Theory of Types, because of difficulties it raised for the Logician program in the foundations of mathematics. This was an early stage in what became a vigorous extensionalist program in Symbolic Logic, with only a few holdouts: the constructivists, and some others such as Carnap, Church, and Fitch.

Now that intensional problems have become not only respectable but central, I think it might be rewarding to dust off the Ramified Theory of Types, and to reformulate it so that it can be compared with contemporary theories.

4. Consequences for Semantic Representation²⁴

The cognitive expectation is that semantics will work out in much the same way as syntax; semantics will be a theory of semantic competence, describing the "knowledge" of the ideal speaker. One consequence of the Knower Paradox is a proof that the ideal speaker's theory will be inconsistent, given this approach to semantics and other plausible assumptions. For on this model of competence, propositions will be recursively generated, like syntactic structures—and the relation between sentences and propositions will be recursive. Thus, if the ideal speaker's theory contains arithmetic, the relation between sentences and propositions will be expressible in the language. The assumption that the ideal speaker's knowledge is logically closed, which seems exactly like other assumptions warranted by this sort of idealization, leads to a contradiction.

This is simply another way of restating the tension between the cognitive goal of graspability to which I alluded in Section 2. The correct reaction, I think, is that the

²³See [Prior 1961] and [Thomason 1982].

²⁴For the most part, this section summarizes points made in [Thomason 1979], which is still unpublished.

metaphor of the ideal speaker needs reappraisal when semantics is taken to be part of the ideal speaker's competence; it seems to me that the relation between theories of the language and speakers is more problematic than the metaphor suggests. In any case, I believe that if the consequences of adhering to the goal of graspability are worked out carefully, it will lead to strong forms of constructivism, or paraconsistent approaches, or other approaches that involve revision of logical foundations. Any assessment of strong forms of cognitivism should take these consequences into account.

As for the issue of semantic representation, my guess is that we are not going to find any theories that make everyone happy. Years of hard work by the best logical talents have not done this with the Liar, and it would be foolish to expect it to happen with the problems that the Knower poses for semantic representations. We can expect, though, that we will develop a variety of technical solutions, and that these will help to clarify the better alternatives. In seeking for *explanatory* solutions, I would hope that those who favor syntactic representations will be aware of the challenge for their approach posed by the modalities, and in general by the linguistic constructions favored by model-theoretic logicians; and equally, that those who favor nonsyntactic representations will be mindful of the need to say something about epistemological questions. This piece of advice is statesmanlike, and therefore boring. I'm sorry for that—but humility really seems to be the proper attitude when attempting to deal with the semantic paradoxes.

5. Consequences for Computer Science

In all cases in which these problems make contact with a nonphilosophical discipline, such as Computer Science, Linguistics, or Psychology, I am uncertain about how close the encounter is or ought to be. In each of these cases it seems that we have reached a pretty rarefied level of the discipline, at so great a distance from implementations, or linguistic evidence, or human experimentation, that the connection with the discipline is stretched near the breaking point. And maybe this is just as well; I would not want to bet on the future of a science if this depended on a definitive solution to the semantic paradoxes. We have to remember that foundational problems are not a sign of an unhealthy discipline; sciences like physics and mathematics have highly problematic foundations, which concern most practicing physicists and mathematicians very little, if at all.

One long-range approach to foundational difficulties, then, is to leave them to a few specialists—usually a small group of technically minded philosophers and reflective scientists—and to treat them as peripheral. The cognitive sciences, as they develop, could conceivably settle into this sort of relation to the paradoxes of representation. To a large extent, this will depend on how able we are to produce solutions that are explanatory and scientifically fruitful.

I don't want to close with the impression that foundational problems of representation are irrelevant to cognitive sciences. These problems certainly do constrain the theories that can be articulated in disciplines such as Computer Science. And there are plausible, well motivated trains of reflection leading from centrally placed material in the cognitive disciplines to these problems of semantic representation. Finally, in the present climate it would be especially unrewarding to draw hard boundaries between Computer Science and Philosophical Logic—though, as I say, one possible line of long-term development might draw these lines so that the job of worrying about an explanatory solution is left to the philosophers.

The point of all this is that I'm uncertain about the long-range prospects for the theory of representation in Computer Science. At this point I wouldn't recommend investing heavily in any particular solution to the paradoxes. And in a situation where different solutions will be pursued, and new ideas can count for so much, it's very difficult to make very reliable long-range predictions. On the other hand, I do have some short-range advice.

The more closely I have become acquainted with the theories of reasoning that are being presently developed in Computer Science, the more urgently I have felt the need for philosophers to become acquainted with these theories. Philosophers, for instance, seem to know little about knowledge representation, whereas Computer Scientists have learned the relevant philosophy. But in the case of the paradoxes, it seems to me that Computer Scientists still have as much to learn from Philosophers as Philosophers have to learn from Computer Scientists. One short-range consequence for Computer Science, then, is that familiarity with the philosophical literature on the paradoxes is important for research on paradox-related issues concerning semantic representation. Fortunately, [Martin 1984] is an excellent written source; and meetings such as this one provide a pleasant way to speed up the interdisciplinary interactions.

ACKNOWLEDGMENTS

I am grateful to Anil Gupta for comments on part of this paper. I would like to dedicate this paper to my teacher Frederic B. Fitch, who labored for many years on the foundations of semantics. Much of his work on the paradoxes is still being rediscovered.

REFERENCES

- N. Asher and H. Kamp. The Knower's Paradox and Representational Theories of Attitudes. *This volume*.
- J. Barwise and J. Perry. *Situations and Attitudes*. MIT Press, 1983.
- P. Benacerraf. Mathematical Truth. *Journal of Philosophy* 70 (1973), pp. 661-679.
- P. Benacerraf. God, the Devil and Gödel. *The Monist* 51 (1967), pp. 9-32.
- T. Burge. Semantical Paradox. *Journal of Philosophy* 76 (1979), pp. 169-98.
- T. Burge. Epistemic Paradox. *Journal of Philosophy* 81 (1984), pp. 5-28.
- G. Chierchia. Nominalization and Montague Grammar: a Semantics without Types for Natural Languages. *Linguistics and Philosophy* 5 (1982), pp. 303-354.
- M. Cresswell. *Structured Meanings*. The MIT Press, 1985.
- H. Curry. The Inconsistency of Certain Formal Logics. *Journal of Symbolic Logic* 7 (1942), pp. 115-117.
- D. Dennett A Cure for the Common Code? In *Brainstorms*, by D. Dennett. Bradford Books, 1978, pp. 90-108.
- J. des Rivières and H. Levesque. The Consistency of Syntactical Treatments of Knowledge. *This volume*.
- R. Flagg. Consistency of Church's Thesis with Epistemic Arithmetic. Abstract, *Journal of Symbolic Logic* 49 (1984), pp. 679-680.
- J. Fodor *The Language of Thought*. Thomas A. Crowell Co., New York, 1975.
- A. Gupta. Truth and Paradox. *Journal of Philosophical Logic* 11 (1982), pp. 1-60.
- A. Gupta. The Meaning of Truth. Unpublished manuscript, University of Illinois at Chicago, 1984.
- H. Herzberger. Naive Semantics and the Liar Paradox. *The Journal of Philosophy* 79 (1982), pp. 479-497.
- M. Kremer. *Logic and Truth*. Ph.D. Dissertation, University of Pittsburgh, 1986.
- S. Kripke. Outline of a Theory of Truth. *Journal of Philosophy* 72 (1975), pp. 690-716.

- M. Löb. Solution of a Problem of Léon Henkin. *Journal of Symbolic Logic* 20 (1955), pp. 115-118.
- R. L. Martin, ed. *Recent Essays on Truth and The Liar Paradox*. Oxford University Press, 1984.
- R. L. Martin, ed. *The Paradox of the Liar*. Yale University Press, 1970. (Second edition with supplementary bibliography, Ridgeview Press, 1978.)
- R. Montague. Syntactical Treatments of Modality, with Corollaries on Reflexion Principles and Finite Axiomatizability. In *Formal Philosophy*, by R. Montague, Yale University Press, 1974.
- L. Morgenstern. A First-Order Theory of Planning, Knowledge, and Action. *This volume*.
- C. Parsons. The Liar Paradox. *Journal of Philosophical Logic* 3 (1974), pp. 381-412.
- D. Perlis. Languages with Self-Reference I: Foundations. *Artificial Intelligence* 25 (1985), pp. 301-322.
- G. Priest. The Logic of Paradox. *Journal of Philosophical Logic* 8 (1979), pp. 219-241.
- A. Prior. On a Family of Paradoxes. *Notre Dame Journal of Formal Logic* 2 (1961), pp. 16-32.
- H. Putnam The Meaning of Meaning. In *Language, Mind and Knowledge*, ed. Keith Gunderson. University of Minnesota Press, 1975, pp. 131-193.
- W. Reinhardt. The Consistency of a Variant of Church's Thesis with an Axiomatic Theory of an Epistemic Notion. Unpublished Manuscript, University of Colorado, 1983.
- J. Reynolds. Types, Abstraction and Parametric Polymorphism. In *Information Processing 83*, ed. R. Mason, Elsevier Science Publishers, 1983, pp. 513-523.
- R. Stalnaker. *Inquiry*. MIT Press, 1984.
- J. Stoy. *Denotational Semantics: the Scott-Strachey Approach to Programming Language Theory*. MIT Press, 1977.
- A. Tarski. The Concept of Truth in Formalized Languages. in *Logic, Semantics, Metamathematics*, by A. Tarski, pp. 152-278, Oxford University Press, 1956.
- R. Thomason. Indirect Discourse is not Quotational. *The Monist* 60 (1977), pp. 340-354.
- R. Thomason. A note on Syntactic Treatments of Modality. *Synthese* 44 (1980), pp. 391-395.
- R. Thomason. Some Limitations to the Psychological Orientation in Semantic Theory. Unpublished manuscript, University of Pittsburgh, 1979.

- R. Thomason. Paradoxes of Intentionality? Unpublished manuscript, University of Pittsburgh, 1982.
- A. Visser. Four Valued Semantics and the Liar. *Journal of Philosophical Logic* 13 (1984), pp. 181-212.
- A. Whitehead & B. Russell. *Principia Mathematica*, first edition, vol 1. Cambridge University Press, 1910.
- P. Woodruff. Paradox, Truth and Logic. Part I: Paradox and Truth. *Journal of Philosophical Logic* 13 (1984), pp. 181-212.