

On Formal Versus Commonsense Semantics

David Israel
AI Center and CSLI
SRI International

There is semantics and, on the other hand, there is *semantics*. And then there is the theory of meaning or content. I shall speak of *pure mathematical* semantics and *real* semantics. I have very little idea what “formal” means in “formal semantics”—unless it simply means semantics done rigorously and systematically.¹ I have even less idea what is meant by “commonsense semantics”. I shall not speak much of the theory of meaning. The distinction between these two modes of semantics, the mathematical and the real, is not meant to be a hard and fast distinction—nor, most assuredly, is it intended to be rigorous or systematic. As I see it, the distinction has primarily to do with purposes or goals, derivatively, with constraints on the tools or conceptual resources considered available to realize those goals. In particular, real semantics is simply pure mathematical semantics with certain goals in mind and thus operating under certain additional constraints. Another way to put the point: some work in pure mathematical semantics is in fact a contribution to real semantics; however, it does not have to be such to make a genuine contribution to pure mathematical semantics.² Hence, since real semantics *can* be executed with the same degree of rigor and systematicity as must all of pure mathematical semantics, it *should* be.

Have I made myself clear? Not *entirely*, perhaps. Let’s try a more systematic approach. Pure mathematical semantics is either a part of or an application of mathematical logic. Real semantics, even though an application of mathematical logic, is a part of the theory of meaning or content. Contributions to real semantics had better cast some light on naturally occurring phenomena within the purview of a theory of meaning—on such properties and relations as truth analyticity, necessity, implication.

Traditionally (indeed, until Montague, almost undeviatingly) the techniques of pure mathematical semantics were deployed for formal or artificial languages. But this by itself is of no importance. These languages were invented, and are of interest only because, or insofar as, they are plausible and illuminating models, in the intuitive sense, of real phenomena in thought and in informal language. Consequently, the nature of the languages studied need not make an essential difference.³ What does make a difference is the purpose

¹This is mildly disingenuous; talk of “formal semantics” is usually grounded in one or another idea of “logical (or, more generally syntactic) form”. But one should beware the overly eager application of such notions to the semantics of natural languages.

²Of course, problems pursued for purely technical, mathematical reasons often turn out to be related to important questions and issues in real semantics.

³Indeed, the two examples I shall consider concern the semantics of formal languages.

or *end* of the study and the fact that the end imposes constraints on the choice of means.

In doing work in semantics, the logician has a range of tools available and certain criteria for choosing among them. In pure mathematical semantics, the only criteria are, in a suitably broad sense, technical. There are no nontechnical constraints; **anything goes**. That is, even if the development of a class of logics was inspired by an attempt to model various pretheoretic notions, work on the pure mathematical semantics of the languages can still diverge quite far from that original motivation. The objective is to provide a systematic way of assigning mathematically describable entities to the nonlogical expressions and mathematically characterizable operations to [or correlating them with] the logical constants so that the best proofs of the strongest and most general results may be achieved. Not so for work in real semantics. There the choice of tools and conceptual resources should be grounded somehow in the nature of the phenomena to be analyzed or, to put it differently, problems in real semantics generate not-purely-technical criteria for choosing among technical means.

This, I realize, is all rather vague and airy—so let's get down to cases. The first illustration is from work on higher order logics, in particular Henkin's proof of completeness in the theory of finite types [1]. The intended interpretation of the relevant language is that the individual variables—those of type 0—range over some domain of elements T_0 , and that for each n , T_{n+1} is the power set of T_n , that is, the set of all subsets of T_n . Monadic predicate variables of type $n+1$ range over all elements of the power set of T_n , m -place predicate variables of that type range over the entire power set of the m^{th} Cartesian product of T_n . The theory of finite types can therefore be regarded as a [perhaps noncumulative] version of a part of impure set theory, that is, it formulates a conception of an "initial segment"—up to rank ω —of the set-theoretic universe over some domain of individuals. Now it is a fairly immediate corollary of Gödel's proof that second-order logic—let alone ω -order logic, which is what we are now concerned with—is incomplete relative to this intended interpretation. Yet Henkin proved completeness for a system of ω -order logic. How?

By ingenious hook and ingenious crook, is how. He introduced a wider class of models (interpretations) according to which the sole requirement was that each T_n be nonempty; otherwise, the interpretation of the T_n 's was arbitrary. In particular, it is not required that each T_n be the power set of the immediately preceding type. This approach made it possible for Henkin to reduce ω -order logic to a many-sorted first-order logic, thereby allowing him to obtain soundness, completeness, compactness, and Löwenheim-Skolem results. Henkin's work was an exercise in pure mathematical semantics. The task before him was to provide a class of models for an axiomatic system in such a way as to provide soundness, completeness, and other results—and to do so in whatever way worked, without any thought being given to the interpretation on which the real significance of the system was based.⁴

Now let's move on to the treatment of languages for propositional modal logics.⁵

⁴Of course, quite independent of Henkin's motivations, it could have worked out that the class of models he focused on *was* indeed of real semantical interest. It just didn't work out that way.

⁵There is an interesting twist as regards motivation in this case. C. I. Lewis, the founding father of modern modal logic, was interested in different conceptions of implication (or the conditional), *not* in differing conceptions of contingency and necessity. Of course, on the conventional view, implication simply *is* validity

Modal logics have been studied as axiomatizations of different conceptions of necessity or possibility—or to put it somewhat differently, as axiomatizations of different conceptions of modal facts. The current standard semantical account is in terms of Kripke structures. For our purposes, let us think of these as ordered pairs $\langle \mathbf{K}, \mathbf{R} \rangle$, with \mathbf{K} a nonempty set and \mathbf{R} a binary relation on \mathbf{K} . (Kripke structures are now usually thought of as triples, the third element of which is a distinguished element of \mathbf{K} . I'll return to this briefly later.)

Roughly speaking, what happens is that the elements of \mathbf{K} are used to index classical propositional models or interpretations—that is, assignments of \mathbf{T} or \mathbf{F} to the sentence letters—and the relation \mathbf{R} , which is correlated with the modal operator, is a relation among such indices. (There may be more than one modal operator in which case there will be more than one binary relation.) Now if one thinks of the models as representing ways the world might be or alternative possibilities (or some such), it is not really such a bizarre exercise to follow Kripke's heuristic; the set \mathbf{K} of indices of models is a set of possible worlds and \mathbf{R} is a relation of accessibility or relative possibility among worlds. This heuristic results in a version of an old idea due to Leibniz: *Necessity is truth in all possible worlds.*

The work on model-theoretic semantics for modal languages and logics using Kripke structures is a bit of pure mathematical semantics that is arguably also a real contribution to real semantics. Moreover, the techniques involved have shown themselves to be widely applicable. Thus, besides work on temporal logics, in which the set \mathbf{K} is understood to be a set of times or time slices and \mathbf{R} the relation, say, of temporal precedence, we have work on provability interpretations in which, for example, \mathbf{K} is the set of consistent recursively axiomatized extensions of Peano arithmetic and $T_1 \mathbf{R} T_2$ if and only if the consistency of T_2 is provable within T_1 . There is also, of course, a good deal of purely technical, mathematical work on the Kripke-style semantics for modal languages. As Kripke asks, "What is wrong with the purely technical pursuit of mathematically natural questions even if the original motivation is philosophical?"

Still, the philosophical questions, questions from metaphysics and the theory of meaning, keep insinuating themselves, as they must. If the work on Kripke structures is to be taken seriously as a piece of real semantics, something must be said about these entities called possible worlds and about the relation between them and the classical models they index.⁶ It will not do simply to say, as we can when doing pure mathematical semantics, that \mathbf{K} is just some nonempty set and \mathbf{R} is just some binary relation on \mathbf{K} , that meets such and such conditions. You've got to put up about possible worlds or shut up. I would argue that when you do put up, the best net result is to postulate a family of structures like

or *necessity* of the conditional.

⁶For instance, the distinction between models and indices is crucial, but that very distinction leaves room for the following possibility; there can be distinct possible worlds which are exactly alike as ways the world might be. That is, one and the same model can be paired with more than one index. Is this just an artifact or is it supposed actually to represent something? If so, what? There are things to be said here, things about representing contingent relations to propositions. Never mind what they are though; the point is that when taking work in the model-theoretic tradition seriously, one has to keep in mind that what is being done is, precisely, *modeling*. One must think seriously about what aspects of the proposed model are merely artifacts and what not.

those to be found in the situation-theoretic universe. But that's an argument for another occasion.

I want to make one more point about propositional modal logics. Oddly enough, structures that yielded models for propositional modal logics had been made available as a result of research in Boolean algebra by Jónsson and Tarski [2]. This work had nothing to do with the issues of necessity and possibility; the research was not in the least concerned with modal facts, nor, in fact, with modal languages. As a result of this work (and thanks to the perspicacity of hindsight) structures for modal languages can be seen to be *[proper] relation algebras*. Proper relation algebras are a special case of Boolean algebras with operators; work on them is directly related to results in universal algebra, the metamathematics of algebra, and category theory. For my purposes, though, the crucial aspect of this work is precisely its austere abstractness and generality. This is work in mathematical semantics at its purest. In this framework, even the set \mathbf{K} rather disappears into the background—to be replaced by binary relations on \mathbf{K} , those being the elements of the algebra.⁷ Once again, the Kripke heuristic is available; it's just farther removed from the mathematical action.⁸

The point to stress is a simple, but an important, one: the “reading” of the set \mathbf{K} as a set of possible worlds, and of \mathbf{R} as a relation of accessibility among possible worlds plays no part in the technical development. That heuristic enters precisely when claims of a real semantical nature are to be made for, or on the basis of, the technical results in pure mathematical semantics. And those claims cannot be extricated from more general issues in the theory of meaning.⁹

Earlier I suggested that I had just about no idea what is meant by *commonsense semantics*. Alas, this too was disingenuous of me. Sad to say, my guess is that most adherents of *commonsense semantics* are convinced that rigorous, systematic accounts of the semantics of natural languages are unattainable. In this regard, Schank and Chomsky are bedfellows, however strange. I know of no good arguments for the Schank-Chomsky view.¹⁰ Rather than canvass the various bad arguments that have been trotted out, let me conclude by citing four crucial sources of confusion that may have led many astray. They all have to do with the scope and limits of semantics.

The first is to think that a semantic account of a natural language has to say everything there is to say about the *meanings* or *interpretations* of expressions of the language, with *meaning* and *interpretation* understood very *broadly and informally*. A theory of the

⁷Very nice work exists on relating Kripke-structures for modal logics to relational algebras—or, more specifically, to modal algebras.

⁸As Professor Scott reminded us all at TINLAP, the prehistory of the model-theoretic semantics of modal logic is quite rich and complex. It starts [more or less] with Tarski's work on topological interpretations of intuitionist logic, continued by Tarski and McKinsey in a more general algebraic setting in which they could draw illuminating connections to one of Lewis's systems (S4). But a more complete telling of this tale deserves both a more proper occasion and a better story-teller.

⁹All these questions arise much more sharply in the case of quantified modal logics. Many of these have been canvassed in an important series of papers by Kit Fine, [3]. In any event, work in quantified modal logic simply has not developed in the robust way as has work in propositional modal logic.

¹⁰That's not to say, though, that the naysayers might not, in the end, be right. There are no guarantees of success in this business.

semantics of a natural language, e.g., English, is not (nor is it intended to be) the whole story of that language, minus its syntax [phonology, morphology, etc.] A semantic account *may be* the whole story about a formal language, minus the account of *its* syntax. But that is because formal languages are studied, *not used*.¹¹ A semantic account of [the declarative fragment of] English should be one that, in a systematic and rigorous manner, relates various facts about [or aspects of] the circumstances in which sentences of English are used to various facts about [or aspects of] the truth conditions of those uses—that is, to various facts about the states of the world those utterances are intended to describe. This is a central aspect of meaning or interpretation—where, again, these are pretheoretical notions—but it does not exhaust the subject.

The phenomenon to be studied is the use of language or, if you like, the phenomena to be studied are kinds of uses of English sentences. Each such use is a complex event with many aspects. Those aspects enter into many different kinds of regularities or systematic connections. There are syntactic facts, morphological facts, phonological facts, facts about the surrounding discourse, facts about the nonlinguistic circumstances of the utterance, facts about the connections between states of mind of the participants, and facts relating such states to various objects, properties, and relations in the environment. These facts are related to one another in a wide variety of ways, some of which are the province of semantics and some not. For instance, any theory of language use had better make room for a distinction between the semantically determined information content of an utterance and the total information imparted by that utterance. The former is not the latter; the latter includes information imparted by the utterance in virtue of the interplay of [broadly speaking] pragmatic factors. In short, acknowledging the possibility of real mathematical semantics for natural languages does not imply acceptance of *semantic imperialism*.

Second, semantics, even construed as part of a theory of language use, is not directly a theory of processing. Any real semantics for natural language should comport with good theories about the cognitive capacities and activities of users of such languages. But no theory of semantics can constitute or be directly a part of such a psychological theory. That a semantic theory does not attempt to be a processing theory, or more generally, a part of a psychological theory, is thus no cause for complaint.

The third point is largely about the scope and limits of lexical semantics. The point is that there are limits. Lexical semantics does not yield an encyclopedia. Any semantic account worth its salt will yield a set of [analogues of] analytic truths, sentences such that the truth of utterances of them is guaranteed by the meanings of the words occurring in the sentence (plus, of course, their modes of combination), together with what might be called “the demonstrative conventions of the language”.¹² Any such semantic account, then, will have to distinguish between analytic truths and *world-knowledge*. Consequently, no such semantic account will say everything there is to say about the objects which are the denotations of lexical items. A brief point about the connection between the current and the previous points is *well worth making*. A good theory of natural language processing will

¹¹Still, we should remember what was said earlier about the purposes for which these languages are devised.

¹²If we have picked out a small set of lexical items as *logical constants*, then those analytic truths will be *logical truths*. Of course, we can make that set as large and as heterogeneous as we want.

have to explain how relevant *world-knowledge* is accessed, and used in processing. *GOOD LUCK to such theorists!* In any event, their job is not the semanticist's job.

Fourth, and last: any real semantics for natural language should be a part of or be accommodable within a general theory of meaning—indeed a general theory of mind and meaning. Nevertheless no theory of the semantics of natural language can itself constitute such a general theory. Return with me now to the example of the mathematical, in particular, model-theoretic semantics of modal languages. As remarked earlier, the classical Lewis systems of modal logic *might* be said to express different conceptions of modality, but they don't express them in the sense that they constitute theories of modal facts. *Nor do Kripke-style, model-theoretic treatments of those logics constitute theories of modality.* The latter constitute ways of thinking about modal facts expressible in the former—that is, they provide *models*, in the intuitive sense, of the phenomena of modality. Kripke, for example, has presented bits and pieces of a *theory* of modal facts in **Naming and Necessity**, a piece which contains no mathematical semantics. David Lewis presents another conception of modal facts in his recent **A Plurality of Worlds**; that book too is devoid of the machinery of model-theoretic semantics. Those different theories may lead to the adoption of different mathematical treatments of modal languages. They will do this precisely by motivating choices among alternative pure mathematical semantic treatments—that is, by providing criteria of choice of a real semantics for modal constructions.¹³

References

- [1] L. Henkin, "Completeness In The Theory Of Types," **Journal of Symbolic Logic**, Vol. 15, No. 1, pp. 81-91 (March 1950).
- [2] B. Jónsson and A. Tarski, "Boolean Algebras With Operators, Part I," **American Journal of Mathematics**, Vol. 73, pp. 891-939 (1951) and "Boolean Algebras With Operators, Part II," **American Journal of Mathematics**, Vol. 74, pp. 127-162 (1952).
- [3] K. Fine, "Model Theory for Modal Logic: Parts I, II, and III," **Journal of Philosophical Logic**, Vol. 7, pp. 125-156 (1978), Vol 7, pp. 277-306 (1978), and Vol. 10, pp. 293-307 (1981).

¹³The research reported in this paper has been made possible by a gift from the System Development Foundation.