The Role of Artificial Languages[*]

Martin Stokhof

1.   Introduction

When one looks into the role of artificial languages in philosophy of language it seems appropriate to start with making a distinction between philosophy of language proper and formal semantics of natural language. Although the distinction between the two disciplines may not always be easy to make since there arguably exist substantial historical and systematic relationships between the two, it nevertheless pays to keep the two apart, at least initially, since the motivation commonly given for the use of artificial languages in philosophy of language is often rather different from the one that drives the use of such languages in semantics.

Of course, this difference in motivation should not blind us for the commonalities that exists between the two disciplines. Philosophy of language and formal semantics have a common history, and arguably also share some of their substance. Philosophy of language is by and large an outgrowth of work in the analytical tradition in philosophy in the first half of the twentieth century. Both ordinary language philosophy, with its emphasis on the description of actual language use, as well as the more logic-oriented and formally inclined school of logical positivism contributed to the definition of philosophy of language as a separate philosophical discipline, with its own set of problems and methods to solve them. Another major contributor to the establishment of philosophy of language as a distinct discipline has been modern linguistics, in particular generative grammar in the tradition of Chomsky that became the dominant paradigm in linguisticsin the fifties and sixties of the previous century. And as it happens, both the generative tradition of Chomsky and analytic philosophy in its formal and less formal guises have been important factors in the development of formal semantics as well. Thus, it should come as no surprise that the two have something in common. That the communalities go beyond a common ancestry, but are reflected in substance and methods as well, will be argued later on.

However, be that as it may, it is still a good idea to keep philosophy of language and formal semantics separate, at least initially, since the role that is assigned to artificial languages and the ways in which these languages are employed in both does differ in a number of respects that are worth keeping in mind. Accordingly we will begin with a brief characterization of the way artificial languages play a role in philosophical investigations (section 2). Then we will look in some detail at their use in formal semantics, focusing on the philosophical presuppositions of the way in which they are employed there (section 3). The differences and resemblances are then the subject of section 4.

2.   Philosophy

In philosophical research on language we mostly find artificial languages being used as notational aides, in much the same way as they are used in, for example, analytic metaphysics, in traditional epistemology, or in philosophy of science. The main goal of this use of artificial languages, such as those of first-order logic, modal logic, and the like, is to clarify a particular argument, or to precisely formulate a conjecture or a thesis. We could call this use of artificial languages that of "easing the conversation", as the purpose of it is to make certain elements in an argument or thesis more readily accessible

---

and thereby more easy to evaluate critically (or more convincing as the case may be). Getting the relative scope of quantificational expressions ('everyone', 'most of us', 'some people',... 'always', 'sometimes', ...) and modal operators ('necessarily', 'obliged', 'possibly', 'allowed' ...), or the exact nesting of Boolean connectives and operators ('and', 'or', 'if ... then ...', ..., 'not', ...), right is often crucial for the understanding, and subsequent evaluation, of an argument, or for the grasp of the precise content of a particular thesis. Of course, in many cases the required disambiguation or explication can be provided without taking recourse to the use of formal notation, simply by explicitly stating, for example, the intended relative scope of two quantificational expressions. But such paraphrases may become somewhat cumbersome and complex, and anyway, the use of notation, apart from being helpful, also has a certain aesthetic appeal. A classical case where the use of notation helps to, if not solve, then at least make clear a particular problem is Aristotle's treatment of temporal modalities in the famous sea battle discussion in *Peri Hermeneias*, which has created a large number of diverging interpretations that in part revolve around the question what are the exact scopal properties that should be assigned to the relevant terms in Aristotle's text. Concretely, the question is whether Aristotle holds that the truth of a proposition about a future event depends on causal determinism (in which case we should construct Aristotle's premise as 'It is necessary that p or it is necessary that not-p'), or whether he merely subscribes in the relevant passages to the law of excluded middle (in which case the premise reads as 'It is necessary that (p or not-p)'.

Although this is certainly a proper use of artificial languages, we should also remark that it is not the most deep use. Be that as it may, what is more troubling is that in many cases authors that use a piece of notation forego the trouble of specifying what it means. And this actually may hamper understanding rather than facilitating getting a certain point across. In many cases authors rely on an implicit understanding on part of their readers of the notation they are using. Often, this may be excused: for instance, when one use the quantifier-notation of first-order logic, it would be somewhat pedantic to explicitly state interpretation clauses for the quantifiers, given that there is an accepted, standard way of interpreting them. However, if one uses the box-operator from modal logic, for example, in a partial formalization of the premise or the conclusion of an argument, such a use really makes sense only if one also indicates which interpretation one assigns to it. With such a wide variety of modal systems around, one really can not rely on a common understanding on part of all of one's readers.

This use of notational devices borrowed from various kinds of formal languages is quite akin to the way in which the use of mathematical notation has become part of our general conversational repertoire. We all employ Arabic numerals and simple arithmetical notation all the time, as if they were part of the natural language we use. This requires no knowledge or expertise beyond a basic level of practical ability. Certainly no acquaintance with anything as sophisticated as, for example, the Peano-axioms, or the Frege-Russell definitions of the natural numbers, is assumed here. This use of arithmetic is really all about easing the conversation, i.e., about making certain practical dealings with quantities and their basic properties more manageable.

Such use of notation borrowed from formal languages can very well be regarded as amounting to an extension of the natural language in which it is incorporated. Cf., the characterization that Wittgenstein gives in Philosophical Investigations, section 18 (Wittgenstein 2009):

> [...] ask yourself whether our language is complete; – whether it was so before the symbolism of chemistry and the notation of the infinitesimal calculus were incorporated in it; for these are, so to speak, suburbs of our language. (And how many houses or streets does it take before a town begins to be a town?) Our language can be seen as an ancient city: a maze of

little streets and squares, of old and new houses, and of houses with additions from various periods; and this surrounded by a multitude of new boroughs with straight regular streets and uniform houses.

For a more restricted set of users logical notation, such as is borrowed from propositional or predicate logic, functions in much the same way. They use the notation to facilitate argumentation and debate, but not much hinges (in general, at least) on their having specific expertise in the logical systems that these languages are part of.

This is different in the context of the natural sciences, where the use of mathematics goes substantially beyond a 'mere' notational use. Here it is the formal systems as such, and not just the formal languages that provide the notational tools, that are being used. For example, the use of Riemann geometry in relativity theory is not a expedient notation, but incorporates substantial claims about the nature of space and time. Hence, this formal system can be used to formulate explanations of phenomena and results that have been observed, and to deduce predictions that can be checked by further observation and experiment. The formal system here is more than a tool, it is part of the substance of the theory in question. Accordingly, here the use of a particular formal system is not just motivated by ease of use or similar concerns, it is subject to empirical validation or falsification, at least in principle. In as much as the formal system is part of the theory any experimental result or observed phenomenon that verifies or falsifies the theory reflects on the formal system by implication.

This seems to mark an important difference between the use of formal languages and systems in philosophy and in the sciences. That does not mean that in philosophy it is only the conversational and clarificatory use that we may encounter. Quite another use, less frequent, but arguably more substantial, that is made of artificial languages in the context of philosophical discussions about language is when it are the properties of artificial languages themselves that are used as premises in some argumentation about natural language. Obviously, such a use is not of artificial languages as such (as is the case in the conversational use noted above), but of *systems,* i.e., of languages with an interpretation (which can be model-theoretic or proof-theoretic). An interesting example of such an application is Putnam's appeal to the Löwenheim-Skolem theorem for first order logic, as part of an argument against the possibility of providing natural languages with a referential semantics. (Cf., Putnam 1983; for a more concrete version of the argument, cf., Putnam 1981.) The theorem states that every countable first-order theory that has an infinite model, has a model of size $k,$ for every infinite cardinal $k.$ This means that such theories are not able to fix their models up to isomorphism. Putnam uses this result to argue that meaning does not fix reference, since given a set of true sentences in a given language we can always by permutation of the objects in the domain change the reference of subsentential terms without this affecting the truth values of the sentences in question. Which shows, according to Putnam, that meaning as such does not fix reference.

It is important to note that an argument like this works on the basis of an assumption: that natural languages are sufficiently like formal languages in the relevant respects so as to allow a transfer of properties and results from formal languages (or systems) to natural languages. In this case, the assumption must be that any natural language 'contains' something like a first-order language and its associated model theoretic semantics. In many respects this use of formal languages in philosophy of language is akin to the use that is made of, e.g., epistemic logic, or probability theory, in formal epistemology. On the assumption that the phenomenon under consideration (natural language meaning, or epistemic reasoning) has a hidden structure that is sufficiently similar to what is explicit in some formal system (first order logic, or epistemic logic) one project results obtained with regard to the latter (such as (un)decidability) onto the former.

One of the problems in this type of use of formal languages is that the background assumption (such as: 'natural languages contain first-order logic') are hardly ever discussed explicitly, let alone justified. But if we really use formal languages in this way, i.e., as formal systems that are part of a theory and that are instrumental in the deduction of certain conclusions, then we rely, not only on the applicability of the system as such, but also on a prior decision to use this system rather than some other one. That priori decision of course calls for independent motivation. In the natural sciences such independent motivation comes from explanatory and/or predictive success with regard to independently acquired experimental and observational results. However, also in this type of philosophical application that is exactly what seems to be lacking.

Yet another use is that is made of artificial languages is when a particular language, or rather system, is appealed to as an arbitrator in order to decide certain issues that are in and of themselves not really linguistic in nature. For example, classical first order logic is often held to encompass an appropriately parsimonious ontology, and thus is used as a standard to determine whether or not some entities that are assumed by some philosophical argument or theory actually are 'proper' entities. Quine's work on ontological commitment provides a classical example. This is a restrictive use, one that intends to single out a limited range of bona fide entities. On the other hand of the spectrum we can locate positions that claim that whatever can be analyzed in some sufficiently general mathematical system, for example Zermelo-Fraenkel set theory (with or without the axiom of choice), or category theory, is something that one might justifiably appeal to. In both kind of uses, however different, it is the formal system that is used as a standard for what can, or can not, be done in and through natural language. However, like in the previous kind of use, this type of argumentation really depends on the availability of external justification for the choice of the particular formal system that one employs as a standard. And like there, such a justification seems to be lacking.

However, one might think that it is formal semantics as the empirical study of natural language meaning that uses similar formal languages, viz., those of logic and mathematics, as its main tools, actually provides such a justification. Therefore, it is now time to look at the use of artificial languages in that realm.

3. Semantics

The use of artificial languages in the study of the meaning of natural languages has started taking off in a systematic manner only at the end of the sixties of the last century. At the time the rise of generative grammar in the Chomsky tradition had done much to expel philosophical doubts as to whether natural languages are systematic enough to allow the application of formal tools, such as those from model theoretic semantics, to them. Tarski had stated in his famous 1944 paper 'The Semantic Conception of Truth' that:

> The problem of the definition of truth obtains a precise meaning and can be solved in a rigorous way only for those languages whose structure has been exactly specified. For other languages – thus, for all natural, 'spoken' languages – the meaning of the problem is more or less vague, and its solution can have only an approximate character. Roughly speaking, the approximation consists in replacing a natural language (or a portion of it in which we are interested) by one whose structure is exactly specified, and which diverges from the given language 'as little as possible.'

For some this meant that in order to become susceptible to any kind of formal treatment natural languages have to be reformed and regimented (beyond recognition, as some others complained). And yet other philosophers took the same 'observation' as

evidence that any attempt to treat natural languages on a par with formal ones was wrong-headed to begin with, and that one would do better to study natural language in an informal, and much more descriptive manner. This approach, pioneered by Austin, Warnock, Ryle, and others, commonly known as 'ordinary language philosophy, was quite aptly also described as 'linguistic phenomenology'.

Both views, however, start from the assumption that natural languages are indeed not systematic enough to allow formal treatment, which is, of course, a complaint that has been leveled against natural languages by philosophers for centuries. The work of Chomsky in generative linguistics apparently inspired much more confidence in philosophers and logicians that perhaps natural languages weren't as unsystematic and misleading as their philosophical predecessors had made them out to be after all.

To be sure, there has already been exceptions, such as Hans Reichenbach, whose Elements of Symbolic Logic (Reichenbach 1947) contains a large section devoted to the application of logic in the description of natural language phenomena, parts of which (in particular his treatment of the natural language tense system) became very influential much later on. And the potential relevance of Chomsky's work for philosophical semantics was already noted in 1953 by Yehoshua Bar-Hillel who wrote a programmatic paper in which he explores the possible connections. This attempt at co-operation met with a negative reaction from Chomsky himself, who throughout his entire career would remain hostile to the idea of some form of model theoretic semantics being of any relevance to linguistics. So it is rather despite this reaction from the leader of the generative movement that at the end of the sixties formal semantics began to flourish nonetheless.

A prominent example is Donald Davidson who claimed, in his seminal paper 'Truth and Meaning' from 1967, that:

> Philosophers of a logical bent have tended to start where the theory was and work out towards the complications of natural language. Contemporary linguists, with an aim that cannot easily be seen to be different, start with the ordinary and work toward a general theory. If either party is successful, there must be a meeting. Recent work by Chomsky and others is doing much to bring the complexities of natural languages within the scope of serious theory.

This sentiment is echoed in many of the early papers that constituted formal semantics as a discipline at the crossroads between linguistics, philosophy and logic: generative linguistics shows how to capture the syntax of natural languages in a systematic and formal theory; in order to extend this to semantics, philosophy provides the necessary conceptual apparatus, consisting of analyses of meaning, reference, and truth that in many respects go back to the early days of analytic philosophy; and logic contributes the formal tools with which these concepts can be applied in a systematic fashion to natural language.

Of course, this co-operation between philosophers, linguists and logicians took on different forms, and not everyone applied formal languages in formal semantics in exactly the same way. One of the most influential approaches turned out to be the one pioneered by Richard Montague. in his 1970 paper 'Universal Grammar', which outlines the theoretical machinery behind what became known as 'Montague grammar', he states:

> There is in my opinion no important theoretical difference between natural languages and the artificial languages of logicians; indeed, I consider it possible to comprehend the syntax and semantics of both kinds of languages within a single natural and mathematically precise theory. On this point I

differ from a number of philosophers, but agree, I believe, with Chomsky and
his associates.

This is a strong statement. No doubt Montague was well-aware of the numerous differences that exist between the formal languages of logic and mathematics and natural languages, yet apparently he is also of the opinion that these are not 'important theoretical differences' that would constitute an obstacle to describing both with the same mathematical means.

Montague's statements suggests that formal languages and natural languages are on a par, but that is not how formal semanticists actually construe the relationship between the two. For their focus of interest is natural language meaning, and hence they use formal languages as tools in their investigations. The most explicit form this takes is when an interpreted formal language is used as a model for (some part of) natural language. The idea behind this methodology is that we can describe and explain the semantic properties of natural language expressions by setting up a systematic relationship between these expressions and those of some suitable formal language, and then use the semantics of the latter to go proxy for that of the former. This procedure of 'indirect interpretation', as it is often called, relies on the assumption that the formal system that is being used in this way somehow provides explanatory power with regard to the natural language. This approach was made popular by Montague's paper 'The Proper Treatment of Quantification in Ordinary English' (Montague 1973), but quite similar ideas can be found on early work of Lewis (1972), Cresswell (1973), and others.

This methodology in essence has remain unchanged from the early days of formal semantics until pretty much the present-day. Like back then, formal semanticists model a particular phenomenon, say, anaphora, or aspect, or vagueness, by either taking an existing formal system that has been developed independently and, in most cases, for different purposes, or by defining one themselves, and then use it to model the natural language phenomenon by providing a more or less strict and systematic translation of the relevant part of natural language into the language of the formal system. In the old days, the tendency was to stick to one particular formal system, that, of course, then had to have all the expressive power one could imagine one would need at some point. Montague used a system of intensional higher-order type theory, Cresswell went all the way and just used ZFC set theory, and Davidson chose to restrict himself, for independent philosophical reasons, to the use of standard first order logic. In particular in linguistic circles, Montague's choice initially won the day, but over the years subsequent developments have seen a plethora of formal systems being used by formal semanticists, unfortunately more often than not without much attention being paid to the overall compatibility of these systems. Domain theory, property theories, belief revision systems, event calculus, different many-valued logics, various non-monotonic logics, dynamic logic, various forms of game theory, second-order type theory, Martin-Löf's type theory, untyped lambda-calculus, Boolean algebras, lattices of various kinds, set theory with or without ur-elements: basically everything in the book has been thrown at natural language phenomena at some point. And then there are the 'custom built' systems, such as various system of discourse representation theory, or situation theory.

What is interesting about the use of this wide variety of formal systems is how the choices are motivated. In many cases, especially those where the semanticists use existing frameworks, they aren't, at least not explicitly. In such cases one rather gets the impression that the formal system that is being employed isn't so much chosen for any particular properties it may have, but simply because 'it gets the job done'. For that to be the case it must of course have enough, and the right, expressive power, but in most cases this is, as it were, settled on incidental, almost ad hoc grounds, not on the basis of a prior investigation of the properties of the system in question. In some cases, though,

the latter do play a role and may actually turn up in the motivation for the use of a particular system. Relevant considerations here usually center around properties such as (un)decidability, (in)completeness with respect to some proof system, expressive power (e.g., first order versus higher order quantification), and similar concerns. Such are, of course, properties of formal systems as such, not necessarily of the natural languages that are studied with the aid of them. However, the questions and concerns that are raised almost always relate to the natural languages, not to the formal systems. Thus, a decision to choose a formal system that is decidable can be found to be motivated by concerns about the learnability of the natural language, or about its effective, practical applicability in, say, common sense reasoning. The assumption behind this is, then, that the relevant properties of the formal system that is used in the analysis of the formal language in some way can be translated back to the natural language, and in that way address the kind of concerns just mentioned.

Thus, artificial languages and formal systems are being used in the analysis of the semantics of natural languages in different ways and with different motivations. However, what appears to be a common assumption is that the formal system is regarded a model of the natural language. In actual analyses and descriptions it is, of course, never considered to be a model of an entire natural language, since it is always only certain aspects of a language (some class of expressions, a particular type of construction) that are under discussion. But in each case what is being analyzed and described is taken to be modeled by the formal system in the sense that all the relevant properties of the natural language are assumed to be adequately represented by properties of the formal system. For it is only on that assumption that it makes sense to think of the formal system as a representation of the relevant semantic properties of the natural language.

This is a widely held view, most of the time simply assumed, sometimes explicitly stated. Textbooks are often the best source of statements of such basic points of view. The following is a quotation from a introduction to natural language semantics by Henriëtte de Swart (De Swart 1998):

> Given that direct interpretation of natural language with respect to the outside world (or some model of it) is not always easy, many semanticists opt for the indirect approach. We know that a translation can sometimes help us to determine the meaning of an expression. Suppose I speak French, but you don't, and we both speak English. In that case, I can teach you something about the meaning of a French expression by translating it into English […] The same 'trick' can be used with the translation of natural language into some formal language. Suppose I can describe the meaning of an English expression by translating it into an expression of a formal language. Because there will be a full and explicit interpretation procedure for the expressions of this formal language, I will immediately have grasped the meaning of the English expression. Of course, I will only have access to it in an indirect way, namely via a translation procedure, but as long as the translation is perfect, the exact meaning will be captured.'

This clearly illustrates that the idea of formal languages as models of natural languages is built right into the core methodology of formal semantics, viz., that of studying natural language by providing translations of relevant fragments into the language of some formal system. The same idea can be found in other textbooks (cf., e.g., Chierchia & McConnell-Ginet (2000): 'Is it possible to regard logical form […] as providing us with a theory of semantic interpretation, with a theory that characterizes what we

grasp in processing a sentence? [...] We think it is possible, as our logical forms do meet the main requirements that semantic representations are generally expected to meet.') However, it relies on two assumptions with regard to its object of study and the way in which this can be accessed.

It is clear that using formal systems to study natural language meaning in this way, i.e., by devising translations between the two, works only if we can assume that both the meanings of the expressions of the formal language as well as those of the natural language are determinate and available prior to the analysis being carried out. They are determinate in the sense that they are able to guide the translation, in the sense of providing the necessary criteria for determining when the translation is actually correct. And they are available in the sense that as semanticists we have access to them independent of and prior to the use to which we put them. (For a more systematic analysis of the determinacy and availability assumptions, including their role in concerns with language reform in the work of Frege, Russell, early Wittgenstein, and others, cf., Stokhof 2007)

Notice that we need these determinacy and availability assumption with regard to both the meanings of the natural language expressions to be analyzed as well the meanings of the formal language that are used. Both are needed for otherwise there would be no way to judge the correctness of the translation, and this is essential if the methodology of indirect interpretation via translation is to succeed. This means that prior to the specification of the translation both are assumed to be given: we need to know what the meanings of the natural language expressions and what meanings are assigned to the formal language expressions before we can start defining the translation or judge the correctness of any attempts. Consequently, what such an indirect specification of meaning does, at best, is representing them in another way than is done by the original natural language expressions themselves. This may very well lead to a more perspicuous representation, or one that has other, technical advantages, but the one thing this methodology will not do is to actually provide meanings for natural language expressions, nor will it allow us to discover their formal properties, simply because prior to the analysis these are assumed to be determined and to be already available.

So where does this leave us with regard to how and why formal languages are used in semantics? For that we need to go back to philosophy.

4.   Semantics and philosophy

The way in which interpreted formal languages are used in formal semantics to model relevant aspects of the semantics of natural language has much in common with the 'Putnam'-type applications in philosophy of language. Properties of the formal system are used to explicate and explain certain feature of natural language meaning, which is made viable to begin with by the assumption of sufficient similarity between the two domains. As we noted above, this calls for independent motivation and justification, which philosophy appears not quite in the position to provide. In the sciences this is different: there data from experiments and observation can be used to test whether the assumed similarity between the formal system and the empirical domain indeed obtains. This works, first of all, because the data used for testing can be obtained independently (at least in principle), and, secondly, because the formal system is integrated in the theory about the empirical domain.

The question now is whether in formal semantics independent justification of these assumption can be procured as well. If so, then formal semantics is like the sciences and it may even function as providing the needed empirical justification for some of the philosophical analyses that make use of a similar methodology. If not, then formal semantics seems to be in the same boat as philosophy.

That the very methodology of formal semantics rests on the assumptions of determinacy and availability is ample reason to think that the required independent and prior justification is lacking here: the method of modeling natural language meaning by means of translation into expressions of an interpreted formal language assumes rather than justifies that the required similarities between the former and the latter exist. So by itself this use of formal systems does not show that it is adequate, and neither do the predictions and explanations that are based on it.

This point is reinforced by the observation that, unlike in the case of the sciences, where the formal system that is being employed is built right into the heart of the theory, in formal semantics no such intrinsic link seems to exists. In view of the wide variety of formal systems that is being used, and taking into account that the choice of formal systems does not appear to be limited by any substantial empirical argument, we must conclude that these systems play a very different role.

What might that role be? It seems we can make the following fundamental distinction between two ways of looking at what it is that we do when we apply artificial languages in the study of natural language. First of all, we can look at a formal language as a model of a natural language. This is the traditional, and still dominant, view on the matter that we have outlined above, which encounters the problems of justification that we also discussed. Second, we may consider a formal language as a tool, –one of the many tools, we should add–, in the study of natural language meaning. We can use a particular formal language to display certain inference patterns involving generalized quantifiers, say, and use another one to deal with, for example, the division between lexical and world knowledge. We can employ the compositionality with which we endow our formal languages to test whether a certain fragment of a natural language allows for a similarly compositional description. (But note that much here depends on what we taken compositionality of natural languages to consist in to begin with, a question that arguably is not straightforwardly empirical.) What is crucially different on this second approach is that the adequacy criteria for our choice and employment of a particular formal language have to come from elsewhere: they derive from the practical concerns that we have, which by themselves are completely agnostic with respect to the tools that we may use or need to use. (And these concerns are even agnostic as to whether logical languages are the best tools. That, then, also becomes a matter that is up for discussion: for some purposes stochastic tools may be arguably better.)

On this second approach we employ formal languages in the study of natural languages not because 'there is no important theoretical difference' between the two, but because they are useful tools. A formal language is not a model of a natural one, but rather a tool that can be used to provide a 'perspicuous representation' of some part or aspect of it. In quite a similar vein one might look at the use of formal languages in philosophical analysis of language: there, too, the primary aim that is served by their employment is perspicuity and clarity. Of course, the content in both is different: formal semantics deals with empirical phenomena, philosophical analysis with conceptual considerations. But the use both make of formal languages as means for perspicuous representation is basically the same.

References

Chierchia, G. & McConnell-Ginet, S. (2000) Meaning and Grammar. Cambridge, Mass.: MIT Press
Cresswell, M. (1973) Logics and Languages, London: Methuen
Davidson, D. (1967) "Truth and Meaning," Synthese 17
Lewis, D. (1972), "General Semantics" Synthese 22:18–67
Montague, R. (1970) "Universal Grammar," Theoria 36: 373-98
Montague, R. (1973) "The Proper Treatment of Quantification in Ordinary English,"

in J. Hintikka, J. Moravcsik, and P. Suppes (eds) Approaches to Natural Language, Dordrecht: Reidel

Putnam, H. (1981) "A Problem About Reference" in Reason, Truth and History, Cambridge: Cambridge University Press

Putnam, H. (1983) "Models and Reality," in Realism and Reason, Cambridge: Cambridge University Press

Reichenbach, H. (1947) Elements of Symbolic Logic, New York: Macmillan Co.

Stokhof, M. (2007) "Hand or Hammer? On Formal and Natural Languages in Semantics," The Journal of Indian Philosophy, 35(5-6): 597-626

Swart, H. de (1998) Introduction to Natural Language Semantics. Stanford: CSLI

Wittgenstein, L. (2009) Philosophische Untersuchungen / Philosophical Investigations, Oxford: Blackwell, 4th revised edition