

Besieging model-theoretic semantics

Radek Ocelák*

January 25, 2016

Abstract

This paper is a philosophical and methodological reflection on the model-theoretic, or formal, approach to natural language semantics. We may distinguish two basic conceptions of the model-theoretic project: a “substantial” and an “instrumental” view. The latter, presenting formal semantics as consisting in building models for the prediction of empirical semantic facts, seems to prevail heavily nowadays. Here, unlike in the “substantial” understanding, the theoretical entities employed in formal semantic modeling need not have any relevance besides theory-internal one. First, I point out a particular problem with the concept of denotation as usually taken for granted in formal semantics. While only the “substantial” understanding of the model-theoretic project is endangered by this problem, I further show that even the “instrumental” way of practicing model-theoretic semantics, seemingly more secure from philosophical critique, suffers from some grave methodological problems. More specifically, I focus on the way lexical meaning is systematically left out of the model-theoretic picture; I argue that formal semantics does not meet the standards of empirical science as concerns evaluation of models against empirical data; I suggest that the standard use of typed lambda-calculus in model-theoretic semantics provides much less semantic understanding than has been hitherto admitted; and I elaborate the worry that an important part of formal semantic accounts consists in merely copying our ordinary ways of talk into the underlying ontology, followed by a more or less trivial redescription.

Keywords: formal semantics, denotation, lexical meaning, evaluation, lambda-abstraction, ontology

1 Introduction

1.1 Model-theoretic semantics

During the 1970s, a new influential paradigm of linguistic semantics was established in the works of people such as David Lewis, David Kaplan, Barbara

*I am grateful to Jaroslav Peregrin, Vít Punčochář, Stefan Kaufmann and a couple of reviewers, who are anonymous to me, for their helpful comments on the previous versions.

Partee, Renate Bartsch, Theo Vennemann, Helmut Schnelle, Max Cresswell, Donald Davidson, and Richard Montague, the most influential (as regards linguistics) of the founders' generation. (Cf. Partee, 1996.) It became widely known as “formal semantics”. There are, however, various ways of being formal, and to a large extent, the tradition of formal semantics can be characterized more specifically as *model-theoretic*. Describing or explicating the semantics of a natural language in a model-theoretic way, in my understanding of the term, involves the following: one needs to show how natural expressions can be adequately translated into some sort of artificial language (say, first or higher order predicate logic); the language in question is to be constituted by an explicit definition of well-formed expressions and provided with a model-theoretic interpretation (in the sense of model theory as a discipline of mathematical logic). That is, objects of various types are to be assigned to the simple expressions of that language as “semantic values”, and this needs to be accompanied with a definition stating how the values of the complex expressions can be retrieved from those of the simple ones by which they are formed. Standardly, a part of this definition consists in determining the conditions under which a sentence of the formal language is true (absolutely or with respect to a given model); therefore *model-theoretic* semantics is to be seen as a species of *truth-conditional* semantics.¹

The present paper is meant as a critical reflection on the classical model-theoretic approach to natural language semantics. One might argue that the inadequacies of the classical approach are nowadays common knowledge, or that by the more recent developments in semantics, the approach has been rendered obsolete, there thus being little need for yet another fundamental critique. I

¹In this introductory characteristic of model-theoretic semantics, I have abstracted over two distinctions that further structure the approach to semantics I am concerned with in this study. First, the issue of direct vs. indirect interpretation: in some cases, the step of translation into a formal language is skipped, and a model-theoretic interpretation is being given directly for the natural expressions in consideration, as in Montague (1970). Second, perhaps more important, the issue of “absolute” semantics: Donald Davidson’s approach to semantics is sometimes (Lepore, 1983) put in opposition to the “truly” model-theoretic stance in that it aims to define the truth conditions of a sentence absolutely, not relativizing to models. However, there have been discussions about whether this is a substantive or merely a superficial difference in semantic practice (cf. Zimmermann, 1999), and these two positions seem to have been coming together lately. Given that “absolutely” can be easily read as “relatively to a single model”, these two options do not make a substantial difference for my purposes.

Let me therefore explicitly state that in this study, I use the attribute *model-theoretic* to refer to most of what has been carried out under the head of formal semantics in the last four decades, as witnessed and sustained by the textbooks Gamut (1991); Heim and Kratzer (1998); Chierchia and McConnell-Ginet (2000), connoting the usual topics, means and methods (such as the extensive use of typed lambda-calculus). The criterion does not consist in the authors’s explicit relativization of truth to a model, but in the use of formal languages with semantics established in the model-theoretic, referential fashion. Some of my points thus also concern the Davidsonian approach to semantics, even though it is set apart from the main flow of formal semantics by a few noticeable characteristics (cf. Chihara, 1975; Clapp, 2002). Although I am aware of the possible confusion caused by my use of *model-theoretic semantics*, I believe the above is sufficient to identify my target, and that less clarification is actually needed than with potential alternatives such as *formal* or *referential semantics*.

would like to oppose such a claim, though.

For one thing, formal semantics in the traditional model-theoretic vein continues to be a vivid discipline within linguistics, and to inspire a good amount of work in the technically conceived sorts of analytic philosophy. Instead of mistrust in the traditional methods, the overall research spirit in these fields seems to be one of steady optimism: altogether, formal semantics has been a success and a crucial advance in the understanding of natural language meaning;² the present analysis of various semantic phenomena may still be imperfect, but with more research effort, we will come to a better knowledge. For another thing, the subsequent developments in formal semantics (say, DRT, situation semantics, or the Amsterdam-style dynamic semantics), while elaborating on particular problematic aspects of the classical approach, have taken many of its basic assumptions aboard. More generally, I believe that the broad logico-philosophical tradition of semantic thinking preserves certain important insights into the nature of human languages. They are insights which the other contemporary perspectives on natural language semantics (such as the lexicographic, the cognitive, or the computational) to their own harm largely ignore (while justly pursuing various aspects of meaning that are neglected within the formal semantic tradition). That makes the mentioned tradition worth studying, even if not worth accepting unconditionally.

For all these reasons, I believe that a thoughtful assessment of the classical model-theoretic approach need not amount to flogging of a dead horse. Without repeating what has been stated in disfavor of that semantic conception on many previous occasions, I hope to present some of its fundamental problems from an original perspective in this study; and I hope the gained insights to be relevant in dealing with the intricate and ever burning question: *whither semantics?*

1.2 Two conceptions of formal semantics

In my opinion, there are two basic ways of understanding model-theoretic semantics:³ it can be interpreted as a “substantial”, or as an “instrumental” theory of meaning. According to the “substantial” understanding, we assign certain entities (such as individual objects, sets thereof, or mathematical functions) to the expressions of our formal language simply because we believe, or hypothesize, that they are the meanings of the corresponding natural expressions (or that something of a similar nature is). On the contrary, the “instrumental” conception does not assume any significance of the employed, often quite abstract theoretical entities, except an internal role in the theory. They need not directly represent anything in reality, but are whatever allows the model to correctly predict certain empirically accessible semantic facts, namely, the relations

²Even a knowledgeable critic of the formal semantic project such as Martin Stokhof begins his paper in an appreciative tone: “Formal semantics is an example of a relatively young, but very successful enterprise.” (2013, p. 205) A more detailed picture of the current “sociological” position of formal semantics and the attitudes towards it that are assumed both within and outside of this research community is provided by Maddirala (2014).

³Also, cf. Stokhof (2002).

of entailment, synonymy, inconsistency and the like between (sentential) expressions, as revealed in the intuitions of native speakers. In this instrumental view, formal semantics is a rather modest enterprise in that it is not supposed to produce heavy claims about the nature of linguistic meaning, or expected to link its own concepts and results with the psychological or neurological evidence about human language processing abilities. (If anything, better fit with this kind of evidence may then be a reason for preferring one model-theoretic analysis over another, equally accurate with respect to the semantic facts to be covered.)⁴

The “substantial” understanding, while dominant in the early decades of logico-philosophical thinking about meaning and arguably still influential in the founding generation of formal semantics (e.g., Cresswell, 1973), seems to play a rather marginal role in the present-day formal semantics. Much has been said and written to the effect that it is hardly a satisfactory conception of the nature of meaning in language. In addition to the obscure character of the assumed relation of denotation between expressions and objects (see below), this conception leaves aside a whole bunch of phenomena that we would intuitively call “semantic” as well. For instance, it has little to offer with respect to the phenomenon of meaning change over time, with respect to acquisition of linguistic meaning and its processing by the human cognitive apparatus in general, or with respect to the phenomenon of semantic competence often being unevenly distributed within a community of (adult) speakers. The contemporary formal semantic research, especially in linguistics proper, can be generally said to proceed in the more humble, instrumental spirit; negotiability of this way of “doing” natural language semantics seems to be consensual.

In what follows, I will first point out an additional problem with the “substantial” understanding, one to which there has not been much attention so far and which, I believe, in itself forces giving up the “substantial” understanding in favor of the weaker, “instrumental” conception. Yet in the rest of the paper, my aim is to challenge even this weaker conception from several angles: I hope to show it suffering from grave methodological problems as well. More specifically, section 3 is focused on lexical meaning and the problem of its connection to the phenomena studied in formal semantics. In section 4, I examine the question of what is the empirical basis for model-theoretic semantics conceived in the “instrumental” way. Section 5 is a brief consideration concerning the role of the function-argument structure that is usually assumed in formal semantic theorizing. Finally, in section 6 I reflect on the apparent problem of transcribing semantic features into the underlying ontology, leading to the possible triviality of the resulting semantic analyses.

⁴We can, however, still contrast this with a purely “engineering” view, which is much more common e.g. in computational linguistics: all that matters from *that* point of view is not the coverage of empirical semantic data, but the success in practical applications such as machine translation. (Cf. Maddirala, 2014, p. 74ff.)

2 Denotation and reference

In the model-theoretic approach, expressions across the board are assigned objects of various types (usually provided by a rich theory of types, in a mathematical sense) as their “semantic values”,⁵ assuming a basic relation of denotation between the expression and the object in question.⁶ I would like to make clear that this relation on the whole is pure theoretical stipulation, with no obvious and with hardly any conceivable counterpart in the reality of meaningful use of natural language expressions.

That there is no observable link between the natural expression “and” (or “at least three boys”) and an abstract object such as a mathematical function (what would it mean for there to be such a link, anyway?) would certainly be too trivial to insist upon, were it not for the fact that *some* expressions (in the first place, proper names) clearly do relate to particular objects in a rather tight way. There need not be any mysterious fine thread drawn between the expression “Barrack Obama” and the current US president: the referential relation in question can be described in use-theoretic terms, even if providing such a description in an exhaustive manner would surely be no easy task.⁷ That is, as concerns proper names, the usual model-theoretic analysis looks very reasonable: a proper name is translated with an individual constant, the constant is interpreted with an individual object of the universe (at any case in an extensional setting, as the simplest case), and the stipulated denoting relation between them reflects an existing referential relation of the proper name in question, albeit in a very simplified way.

Proceeding from proper names to other vocabulary, however, the denoting relation assumed turns from intuitive to highly arbitrary.⁸ The change can be viewed in gradual fashion. It is still fairly intuitive (that is, use-theoretically justifiable) to relate a common name (such as „horse“) with certain objects in the world;⁹ the same is somewhat less intuitive for one-argument verbs (such as

⁵With the exception of expressions that are semantically introduced in a syncategorematic way, such as quantifiers in the first order predicate logic. In more elaborate semantic models, syncategorematic definitions are usually avoided.

⁶Cf. Bach (1989, p. 4): “The [Chinese] word *airen* means the same thing as what is meant by spouse. [...] *Airen* means this thing and *spouse* means this thing. What is the thing that these two words mean or designate? Let us call this thing, whatever it is, the *denotation* of the linguistic expression in question. [...] First, we need to show how to assign denotations to all the basic or lexical elements in the language, Chinese, English, or whatever. And then, we need to show how to put together the denotations of the simple expressions, words like *spouse* or *airen*, and to show how the denotations of complex expressions can be made from the denotations of the simple ones.”

⁷Cf. Groenendijk and Stokhof (2005, p. 101): “[W]hen we focus on the meanings of those attitude expressions that are closely connected with meaning and semantics, such as ‘to intend’, ‘to mean’, ‘to understand’, ‘to refer’, ‘to describe’, and take seriously that the meanings of these very expressions originate, not in soliloquy or through introspection, but in conversation and interactive learning, it becomes clear that the capacities to which these terms refer are social through and through.”

⁸This is noted by Glanzberg (2009, p. 287): “[S]emantics assign semantic values much more widely than the intuitive notion of reference suggests.”

⁹Marconi (1997, p. 104): “First of all, I agree: ‘cat’ refers to cats, ‘spoon’ refers to spoons,

„to sleep“). In any case, we might not want to insist that the common name or the verb refers to *the set* of the appropriate objects, rather than simply to any of them, and even less that it refers to a *mathematical function* that assigns truth values to individual objects. Unary logical predicates are, though, typically interpreted in this way, and some sort of function is what the translations for most of the other natural expressions are typically taken to denote.

Yet *no ordinary natural expression* refers to a mathematical function in a way comparable to that in which the natural expression “Barrack Obama” *really* refers to the current US president – only fairly advanced expressions such as “the function XY” do. (To put it differently, the expression that in this sense *really* refers to the function assigning the truth value 1 to all cats and 0 to other individuals is quite clearly not “cat”, but – of course! – the expression “the function assigning truth value 1 to all cats and 0 to other individuals” and the like.) As obvious as this point is, I believe that it deserves some emphasis, for the attention it gets in the textbooks and applications of model-theoretic semantics is virtually zero. There clearly is a dividing line (maybe not too sharp) between, on one hand, the expressions for which the model-theoretic analysis in terms of denoting an object involves a considerable amount of realism, and, on the other hand, the expressions for which the assumed denoting relation does not stand for anything in the reality of natural language use. In model-theoretic expositions, however, this distinction gets regularly swept under the rug of “denotation”. I do not say that the assigned *object* cannot reflect anything of the reality of use: *e.g.*, the truth function with which logical conjunction (the most basic translation of the natural “and”) is normally interpreted can certainly receive a plausible use-theoretic interpretation. My point is merely that the natural expression “and” does not refer to anything in the way in which a proper name does. (Note that simply insisting on a metaphysics where the denotation relation can relate some expressions to concrete objects like Barrack Obama and other expressions to abstracts such as mathematical functions will not help us out. For then we still have the absurd synonymy between “cat” and “the function assigning truth value 1 to all cats and 0 to other individuals”.)

That, I believe, presents a problem for model-theoretic semantics conceived as a substantial theory of meaning: one who maintains this conception should not pass in silence the dividing line I have just sketched. How come that such an important distinction is ignored in the official formulations of the program? Proper names and the like lend some initial credibility to the general idea that a natural expression can be semantically analyzed assuming a denoting relation between it and another object, but this idea seems to be drawn way too far in model-theoretic semantics. In a sense, this is just the old Wittgensteinian re-

and so forth. That is to say, most speakers regularly use ‘spoon’ to refer to objects of a certain shape and use, to spoons, in short. [...] What is it for a speaker to use the word ‘spoon’ to refer to spoons [...]? Typically, it is to come back with a spoon in one’s hand when one is asked to fetch a spoon; [...] it is to answer questions such as ‘Are there any spoons in the drawer?’ on the basis of whether or not there are spoons in the drawer; and so forth. [...] There is nothing mysterious in the notion of using a word to refer to xs: it is just shorthand for a variety of performances [...] Do they have anything in common? Yes, all of them are concerned with both spoons [...] and the word ‘spoon’.

proach of exercising the “Augustinian picture” of meaning (Wittgenstein, 1967); but in so far as some people keep conforming to that picture, or to the substantial understanding of the model-theoretic approach, it seems worth repeating.

An obvious way out consists in holding that there need not be anything in reality that would directly correspond to the denotation relation; in assuming the position that this relation is nothing but a theoretical device that provides for success on the level of whole theories. That would lead us to the instrumental understanding, and indeed, if we do subscribe to such a conception, the objection above need not be of any concern to us.¹⁰ But there is a price. If the – perhaps not at all comprehensible – question of what reality is *really* like is traded for the question of whether our theories deal with it successfully, then there is this other question that is going to stick with us. Do they?

3 Lexical meaning

A problem with model-theoretic semantics, one which arguably concerns not only the substantial, but also the instrumental understanding of the discipline, is that this approach has little to say on the matter of lexical meaning, or the semantic difference in pairs of lexical units such as *dog* and *cat*, or *brave* and *bold*. A usual qualification is that formal semantics primarily, or even exclusively, pursues the issues of structural or grammatical meaning (that is, *types* of meaning an expression can have, and the ways of their combination), as opposed to the issues of lexical semantics. However, can a semanticist really waive the responsibility for lexical semantic matters in this way, without this undermining the sense of the model-theoretic project as such?¹¹

Take the sentence “Alfred loves Beth”, and its translation in a first order language, $L(a,b)$, assuming that the proper names refer unambiguously. I will just discuss this elementary extensional analysis, as I believe that in the more advanced, intensional semantic models, the problem is the same in principle. Now, L in this translation is a mere letter until we equip it with a denotation. Typically, it is assumed to be interpreted with a set of ordered pairs of individuals such that the first individual loves the second. That, however, does not tell much to anyone who does not know the meaning of the verb *to love*; and the extension of the set is usually not given by any other description. In principle

¹⁰For those not completely convinced about instrumentalism, the least committing version of it may very well consist in nothing over and above the decision that the objection *should* be considered inconsequential.

¹¹Note that parts of the classical formal semantic agenda do concern semantic differences within certain lexical classes (such as quantifiers or indexical adverbs); thus they no doubt warrant the label “lexical semantics” in a looser sense. In a talk given in Tübingen, 2014, Michael Glanzberg even quoted James Higginbotham as saying that model-theoretic semantics makes a particularly good job as “lexicography of logical constants”; an observation I tend to agree with. Also, it should be noted that there are more “grammatical” aspects of the meaning of words (genericity, countability, aspect and the like), which are in fact often in the focus of formal semanticists. In this section, I use the term “lexical semantics” rather narrowly, referring to the semantic differences within open-ended lexical classes (such as, between *dog* and *cat* or *brave* and *bold*).

it *could* be defined without mentioning love at all, saying *e.g.* that the set denoted by “love” consists of the ordered pairs ⟨Quasimodo, Esmeralda⟩, ⟨Romeo, Juliet⟩, ⟨Juliet, Romeo⟩, and so forth – but hardly ever it is. For most lexical units, only the *type* of the assumed semantic value is effectively specified in model-theoretic semantics, not any particular instance of that type. (Cf. Kamp and Stokhof, 2008, p. 60.)

That is, in providing truth conditions of sentences, model-theoretic analysis presupposes knowledge of the lexical meanings in question, without describing them in a non-trivial way. “Alfred loves Beth” is true if and only if the ordered pair ⟨Alfred, Beth⟩ belongs to the set of ordered pairs denoted by *L*; but only someone who already knows the meaning of *to love* is in a position to decide whether the condition on the right holds.¹²

A formal semanticist might want to argue that this is in perfect order: it is not a matter of semantics to discover who is in love with whom, but of empirical science in general. “Let us assume we have three things in our interpretation now: *A*, *B*, and *C*, and somehow we find the set of pairs that are in the denotation of *Love*. How we do that is not part of semantics, but somehow we know that *A* loves *B*.” (Bach, 1989, p. 12.) “We must be very clear about one thing – that we are engaged in logic for linguistics, but we are not engaged in science in general. It is not part of our job to decide when particular statements are true or false about the real world. [...] We are concerned here with the form of truth, not with how we find out whether something is true. [...] We are concerned with the structure of an interpretation rather than the basis for it.” (*ibid.*, p. 24)

But an empirical scientist cannot investigate who is in love with whom¹³ unless she knows what *love* means. If she does not, who if not a semanticist should be the person to ask?! As semanticists, we surely cannot get away with simply saying “well, *love* has in its denotation all the ordered pairs of individuals where the first individual can be said to love the second”. For it is no one but us who is accountable for knowing whether – in the extensional setting, still, but the argument can be made as well for the intensional – the latter *can be said* about a particular pair of individuals.

Sure, one can still resort to the claim that model-theoretic semantics simply focuses on the structural, not the lexical, aspects of natural language meaning. As noted by Marconi (1997, p. 1), though, people in the everyday use preferentially refer with “meaning” to lexical issues: hardly anyone uses that word to point, say, to the problem that the meanings of a determiner such as *at least 3* and of a common name such as *boy* need to combine into the meaning of a nominal phrase such as *at least three boys*. It is therefore at least surprising that the lexical aspect should be completely left out of a semantic discipline; the more so that the discipline in question is not defined by the goal – allegedly,

¹²Cf. Marconi, 1997, p. 109. In effect, this is not getting too far from the Tarskian idea of semantic analysis via disquotation, stating that the sentence “*p*” is true if and only if *p*. As noted by Michael Glanzberg (same occasion), even if this observation is not trivial, it is still terribly boring for a practicing semanticist.

¹³Or who is in a genetic relation to whom, for that matter. Of course we don’t need scientists to tell us of love.

to account for the structural aspects of meaning while avoiding the lexical – but rather by the employed method (that is, by being *formal* and *model-theoretic*). One may wonder, as Stokhof (2002, 2008) does, whether “lexical vs. structural” is a theoretically neutral distinction after all, rather than one constituted by the achievements and failures of a particular (namely, logical) method.

Still, there is nothing wrong with dividing labor – as far as there is some way of putting the pieces back together; that is, of combining a model-theoretic description of the “structural” with an adequate lexical semantics into a functioning whole. Is there such a way? What model-theoretic semantics can offer in the way of accounting for lexical semantic features is the method of meaning postulates (Carnap, 1952, Partee, 1996, p. 34, Marconi, 1997, p. 12ff.); that is, of supplying the model-theoretic analysis of “Alfred loves Beth” with semantic statements such as “if x loves y , then x would do everything in the world for y ”, etc. However, Marconi (1997, p. 18ff.) shows quite clearly that meaning postulates are at best capable of capturing the *inferential* side of the lexical semantic competence (such as, the knowledge that if x is a dog, then x most likely has four legs and a tail). On the contrary, they necessarily fail to account for its *referential* aspects (that is, the ability to correctly apply terms in the world, e.g. the knowledge that *this* is a dog).¹⁴

At the same time, it is hard to find a more adequate method of “doing” lexical semantics which would easily fit with the model-theoretic approach.¹⁵ And as long as there is no good and compatible lexical semantics alongside the model-theoretic enterprise, the latter remains subject to a burning objection. Isn’t one of the purposes of the distinction between lexical and structural semantics that the idea of the former (missing, so far) serve as a wadding to protect the latter (the actually elaborated one) from possible empirical refutation? Such, at least, has often been the case in linguistics with the more general distinction between pragmatics and semantics, or performance and competence. Perhaps, practicing model-theoretic semantics is like assembling a splendid rear part of a time travel machine, regardless of whether a compatible front is being worked upon. As regards our part of the whole damn task, we’re doing pretty fine,

¹⁴There is of course the option of insisting that those referential aspects are not properly a matter of semantics. That however would look a bit like reducing all engineering to driving screws just because we happen to own a screwdriver – presumably the ability to apply an expression in the world is part and parcel of our meaning knowledge. A pertinent question then is: can the fact that *this* particular thing in my garden is a dog be accounted for by any general theory at all? The only answer I can offer is a hesitating one: perhaps we should be then looking not so much for a theory to state so, but rather for a device to apply the term correctly...

¹⁵Baroni et al. (2014); Erk (2014) are attempts to supply the model-theoretic account of semantic composition with concrete lexical (namely, distributional) semantics. While the effort to provide a computationally viable and practically testable semantic account is admirable and rather promising on the lexical part, the whole edifice in my opinion suffers from undue faith in the model-theoretic compositional principles. Given the other criticism formulated in the present study, the proposed combination of distributional and model-theoretic semantics appears to be a matter of will rather than of natural compatibility (as manifested e.g. in the fact that the concept of episodic denotation plays a crucial part in the latter, but hardly any in the former).

aren't we...

4 The empirical basis of semantics

The essence of the instrumental understanding is that the model-theoretic semantic analysis should simply account for the semantic facts concerning the involved expressions, disregarding the question of the realism of the principles or entities employed. *Account for* is a slippery expression in science; here it should be arguably read in the sense of *descriptively cover* or *predict*.¹⁶ The ultimate semantic facts to be covered are usually taken to consist in the truth (acceptability) conditions of sentences and in the relations of entailment, synonymy and incompatibility between them, both as accessible via intuitions of native speakers. (Cf. Partee, 1996, p. 17; Peregrin, 2001).

To the extent of this description, one might think that formal semantics is a standard case of empirical science: here's a model, and here's the empirical data to evaluate the model against, so as to see whether it outperforms an older model; if it does, we may (under certain additional conditions) want to call it progress. That seems to be the situation, e.g., in various branches of computational linguistics, but with formal, model-theoretic semantics it is not even nearly the case.¹⁷ The main problem here is clearly “that great difficulties arise when determining just what the data is that you are trying to account for. In general, we deal in linguistics with native speakers' judgments [...] The data can get very messy. The kinds of judgments that we have to ask about in semantics often seem to be very shaky: How many ways ambiguous is a certain sentence with several quantified noun-phrases in it? Does a certain sentence entail another sentence?” (Bach, 1989, p. 128–129).

I do grant that the problem of defining an appropriate empirical basis for natural language semantics is highly intricate. Yet it is simply intolerable that for more than four decades, as of now, formal semanticists have been steadily formulating further models (that is, analyses of various language phenomena), instead of trying to subject any of them to a proper evaluation.¹⁸

The evaluation part of the scientific scheme has been – and largely still is – replaced in formal semantics by appeal to the authors' *own* semantic intuitions; that is by their subjective sense (maybe followed by the consent of the research

¹⁶Not necessarily in the sense of *explain*, since in various contexts we may demand yet a bit more to call something a proper *explanation* – e.g., an interesting reduction of the phenomenon in question to other phenomena, or anchoring it to the known psycho- and neurological underpinnings.

¹⁷Cf. Maddirala, 2014, p. 74ff. This is however not to put computational linguistics in opposition with all theory whatsoever. To some extent, computational models generally do depend on linguistic theory, both in their construction and in the evaluation against linguistically annotated data.

¹⁸Cf. Partee (1996, p. 36): “Many of the most fundamental foundational issues in formal semantics (and in semantics as a whole) remain open questions, and there may be even less work going on on them now than there was in the seventies; perhaps this is because there is more work by linguists and less by philosophers, so the empirical linguistic questions get most of the attention now.”

community) that the truth conditions and semantic relations captured by their analysis are the right ones.¹⁹ Let me say a few things about the (un)acceptability of such a methodology.

First, the appeal in semantics to the intuitions of any particular speaker involves some non-trivial assumptions, such as the individualism about semantics (the assumption that semantic structures of a language in general are more or less directly reflected in the competence of an adult native speaker) or the accessibility of competence (the assumption that the individual semantic competence does not seriously diverge from what the individual is able to report about it) – cf. Stokhof (2008, p. 236), Stokhof (2007). I want to leave these issues aside here, though – in the end it does not seem that implausible that an individual speaker may still provide us with a fairly useful picture of the overall semantics of the language in question.

Stokhof (2007, p. 621; see also 2011) further raises the following principled objection against taking intuitions for the empirical basis of semantics. “One crucial feature of that account [of linguistics as a descriptive and explanatory theory of linguistic competence] is the concept of tacit knowledge of language. This concept serves two purposes at the same time, and that is where the problems start. First of all, tacit knowledge of language is what the theory explains, since this tacit knowledge [...] is what users of the language employ in production and reception. Second, it is also what constitutes the data of linguistic description and theorizing. When the linguist appeals to his intuitions [...] about semantic properties and relations, such as ambiguity, entailment and the like, he employs that very same knowledge. But this is deeply problematic, for in effect it says that the phenomena we want to describe and explain, and the data on which we base our descriptions and explanations are really one and the same thing. [...] How is this different from other empirical disciplines? [...] In natural sciences the object studied, physical and biological nature, provides data, but does not equal data.”

Technically, one could object that the usual methodology does *not* identify the object of study and the data, since semantic intuitions constitute the object, whereas the data are whatever we manage to record of these intuitions. However, Stokhof at least seems to be right in his unwillingness to conceive the very object of semantics as consisting of the semantic intuitions – for they are precisely intuitions *about* semantic facts, and do not *amount* to these facts.

I suggest that we can save a part of the “intuitionist” methodology of formal semantics by means of a substantive reinterpretation: Semantic intuitions are not semantic facts, nor are they the only or the primary way of accessing these facts. The semantic facts such as synonymy or entailment between expressions can be understood in use-theoretic terms (*e.g.* as elaborated in Brandom, 1998), and the ways of acting that ultimately constitute such relations are publicly observable. That is, we can in principle learn about semantic facts without

¹⁹van Benthem (2013, p. 19): “A few decades ago, the relevant facts of natural language were produced by the same theorists whose systems were at stake, in the form of judgments of well-formedness, possible readings, or supported inferences.” One doubts whether the situation in model-theoretic semantics has changed all that much since those early days.

consulting anyone’s intuitions at all. But at the same time, there is no denying that asking for intuitions is a highly efficient heuristics. Semantics in this respect is no different from, say, geography.²⁰ Surely the geography of a country is not constituted by the intuitions or knowledge of the inhabitants; yet if you are supposed to learn about that country and your budget is severely limited, then going and asking some of the folks for their opinion, instead of performing your own landscape research, is not a choice that miserable.

My ultimate objection to measuring a semantic analysis by the author’s semantic intuitions is rather earthbound: such an author is clearly in a conflict of interest, and the same can be claimed about the whole research community when its relations to the outside world are taken into account. In civilized societies, however, the very emergence of a conflict of interests is a reason to quit some of the concurrent activities. In formal semantics, I am not aware of any established criterion of adequacy other than the judgment of those palpably interested in the impression of overall success, and that is truly alarming. For all we know, the existing model-theoretic analyses may be excellent, moderate, or disastrous. In lack of a respectable evaluation methodology, who can really tell?

Also, little validation seems to be coming from the pragmatic side. When directly asked about practical applications of formal semantic theories, Bach admits that he is not aware of many. “I do not feel the need to justify the pursuit of knowledge of any kind you might mention by the possible practical applications it might have [...] Here is a phenomenon [...] so I want to understand it.” (1989, p. 127) I certainly do not claim that any understanding worth that name must earn you a million dollars, nor that the *only* conceivable validation is pragmatic. Yet the understanding of a natural phenomenon usually opens *some* new interesting possibilities for us. If there is a lack of those, are we all that sure that there is any actual understanding?

True, nowadays it would be blatantly false to claim that author introspection keeps being the only method to assess a model-theoretic analysis of a particular language phenomenon. For instance, in the context of the recent formal semantic research on gradable adjectives, several types of empirical data have been made use of: on one hand, psycho- and neurolinguistic data, on the other hand, semantic judgments elicited directly or in variously indirect language tasks.²¹ However, I am not sure that this has brought the field much closer to the ideal of a transparent evaluation of semantic models.

As to the former sort of empirical data, it is not clear at all what relevance they should have for model-theoretic semantics in the instrumental understanding, with its absence of ambitions in the way of cognitive adequacy (see e.g. Bach, 1989, p. 120; Stokhof, 2008, p. 225; Groenendijk and Stokhof, 2005, p. 103; Marconi, 1997, p. 108). Sure, they can help decide between two equally successful models, but even that is a certain “extra”, something we are officially

²⁰I owe this simile to Jarda Peregrin.

²¹Also, there are corpus-based studies examining the co-occurrences of gradable adjectives. But it is notoriously difficult to distill semantics from a corpus, and what I say in the following about the method of semantic judgments basically applies to the co-occurrence method as well.

not looking for when engaged in formal semantics.

The latter type of data, empirical data on the semantic judgments of subjects other than the author him- or herself, arguably are a sort of evidence (in the sense of the qualification above) against which an instrumentally conceived model can be evaluated, and maybe this is the only practical option. But as noted by Bach, “the data can get very messy”. 57 % of respondents find a certain inference plausible, 24 % do not, 19 % are not sure – that is a common picture with this sort of studies.²² Such data have been used as clues and heuristics in semantic modeling. But I am not aware of any attempts to set up an independent empirical basis which would provide for a general and reliable evaluation of a semantic model’s efficiency, in comparison with competing models of the same phenomenon (cf. Maddirala, 2014, p. 62). Until something along these lines takes place, I am afraid that model-theoretic semantics cannot be considered a standard discipline of empirical science.

5 A note on lambdas

Bach (1989) and Partee (1996) argue that an important contribution of Montague grammar was that it established a tight link between syntax and semantics. “There is in that theory a functional relationship between the syntactic categories and the semantic types of his intentional [*sic*] logic and hence the types of objects in the model structure.” (Bach, 1989, p. 124.) “Another important legacy of Montague’s work [...] is the idea of seeing function-argument structure as the basic semantic glue by which meanings are combined.” (Partee, 1996, p. 21.) “Dependence on first-order logic had made it impossible for linguists to imagine giving an explicit semantic interpretation for “the” or “a” or “every” or “no” that didn’t require a great deal of structural decomposition into formulas with quantifiers and connectives, more or less the translations one finds in logic textbooks.” (Partee, 1996, p. 19-20.) “[L]ambdas provide a particularly perspicuous tool for representing and working with function-argument structures explicitly and compositionally.” (Partee, 1996, p. 24) “[N]atural language syntax suddenly looked much less crazy; instead of the great mystery of how English syntactic structure related to its putative logical form [...] there suddenly arose the possibility that surface structure or something close to it [...] might be very well designed as a logical form for expressing what natural languages express.” (Partee, 1996, p. 24-25.) Let me add a comment on the nature of this formal semantic breakthrough (which made Barbara Partee remark at a certain point that lambdas changed her life; Partee, 1996, p. 24).

What do lambdas add to the classical Russel-style analysis of the sentence “all men are mortal” by means of $\forall x(Man(x) \rightarrow Mortal(x))$? They enable us to provide a separate translation for whatever we regard as a syntactic constituent of the sentence, and to assign each of those syntactic bits its own semantic value so that these values can combine into the meaning of the whole sentence (truth

²²Cf. Stokhof (2008, p. 215) on the problem of interpreting such heterogeneous results in the framework of methodological individualism.

value 1, assuming the extensional setting again) in a compositional fashion. The quantifier *all men* can be now translated with $\lambda X\forall x(Man(x) \rightarrow X(x))$ and interpreted with an object of the type $\langle\langle e, t \rangle, t\rangle$ – that is, with a function such that it takes a function from individuals to truth values (that is, the semantic value of a logical predicate such as *Mortal*) and assigns a truth value to it.

The problem lies in the specification *function such that...* and is quite closely related to my previous worries regarding the position of lexical meaning in model-theoretic semantics. The semantic value for $\lambda X\forall x(Man(x) \rightarrow X(x))$ is actually not given in a very tangible manner. It does not really allow us to say what truth values this function exactly assigns to which functions. To be sure, there *is* a piece of actual semantic information involved. What we get for the determiner *all* is indeed a proper semantic analysis: “All X are Y” is true if and only if the set denoted by X falls within the set denoted by Y. But as neither X nor Y are spelled out, also the more complex functions building on them cannot be more than vaguely outlined, with reference to our everyday lexical knowledge. The denotation of *all men* is only specified as the function such that it assigns to predicate functions the *right* truth values (referring to our intuition about subsumption of sets; e.g., it assigns the value 1 to *mortal* and 0 to *dark-haired*). And how does that differ from merely saying that that the semantic value of *all men* is *whatever* does the job in the right way?

To better focus on the problem, we may consider a case not complicated by the presence of non-trivial semantic information (as in the case of the determiner *all*). The generalized quantifier analysis of proper names would presumably translate Barack Obama’s name with $\lambda X.X(obama)$ and interpret the result with a function assigning truth values to predicate functions. Which function is it, then? Well, the function that assigns 1 exactly to those predicate functions which assign 1 to the individual named Obama (that is, the function assigning 1 to the denotations of the predicates which are true about Obama). But given the lack of lexical determination, this amounts to nothing but saying that the denotation of *Obama* is the function that we need to get the truth values right...

Call it *function-argument structure* as we may, it seems in fact rather trivial, although omnipresent in formal semantics. Having decided that the meaning of the whole is to be compositionally derivable from the meanings of the parts, and having certain ideas about the meaning of the whole and the meaning of the part A, we stipulate that the meaning of the part B is *whatever* gives the former when applied to the latter. (As if we were to design an engine and got satisfied with characteristics such as “Part A: Whatever transmits the combustion energy to the rotation of shaft”.) Where exactly is the breakthrough that lambdas are supposed to have caused in semantics?²³ One could try to argue that the point of the instrumental conception of formal semantics is precisely that the lambda-phrased solution, as conceptually trivial and cognitively inadequate as it perhaps is, empirically *works*. But my reply to that would be that I don’t see how a thing that is basically defined as “anything that works in the right way”

²³Confront this with Glanzberg’s (2009, p. 283) statement that “working out that the type of quantified noun phrases is $\langle\langle e, t \rangle, t\rangle$ amounted to a major advance”.

could possibly not work (or, rather, “work”).

Jointly, the three interlinked problems discussed in the last three sections (that is, lexical semantics, empirical evaluation and the common use of lambda-abstraction) raise a serious question: how justified really is the status of model-theoretic semantics as a superior account of semantic composition?

6 Transcribing semantics into ontology; trivial redescriptions

One more, not the least important worry I have regarding model-theoretic semantics is with another seemingly all-permeating practice, about the admissibility of which there has not been much explicit discussion (but see Peregrin, 1995): copying semantic distinctions into the assumed ontology, followed by a trivial “explanation” or redescription of the former in terms of the latter. I believe this general principle can be illustrated in a variety of formal semantic domains.

6.1 Possible worlds (and times)

Take a model-theoretic classic, the elementary possible world semantics for the modal vocabulary. It is possible that P / possibly P / maybe P / it might be the case that P , if and only if in some of the possible worlds, P . Or: “What do we need to say about when “Necessarily F ” is true in terms of when F is true? The answer is that “Necessarily F ” is true if and only if in every possible world F is true.” (Bach, 1989, p. 26) That is, possible worlds are introduced into the ontology we are employing, in order to explain, describe, or *explicate* the functioning of the expressions such as *maybe*, *might*, *possibly*, *necessarily*, *must* etc.²⁴ Now, most formal semanticists are not modal realists such as Lewis (2001, p. 84 ff.) but modal “moderates”: a non-actual possible world is just a way in which things might be different from what they are (Bach, 1989, p. 27; Stalnaker, 1976). Or, they assume a purely instrumental position concerning the status of possible worlds: “Any set of objects and any [accessibility] relation with the right properties would do just as well” (Gamut, 1991, p. 72).

So, in the moderate position, “it might rain” is true iff in some of the ways in which the world could be it rains. But what is the right side of this biconditional if not a gawky expression of the fact that it might rain?! Unless we believe in some sort of hard reality of possible worlds, the “explication” in question amounts to nothing but a straightforward paraphrase using an artificial term in whose definition we heavily depend on the vocabulary in need of explication. As emphasized by Kripke (1972), the possible worlds we are bringing in our ontology are no distant planets whose particular features we can discover with a good telescope. The only way to argue and come to know anything about

²⁴For the sake of simplicity, I only consider the alethic interpretation of this vocabulary. For the deontic or the epistemic, the following argument can be made as well.

them is based on our standard and appropriate use of the words *might*, *could*, *perhaps*, *possibly*, *must* and the like. Why then would we think that the talk about possible worlds can illuminate the very same terminology in any useful way? “*In every possible world* is nothing over and above *necessarily*.” (Peregrin, 1995, p. 207)

I believe the objection above cannot be avoided by the instrumental position of Gamut (1991), either. Possible worlds may be just whatever helps to account for the semantic behavior of the modal vocabulary. But one can hardly set up such a helpful thing – “a set of objects and an accessibility relation with the right properties”, where the “right properties” must be supposed to characterize both the relation and the objects in that set – without investing a good amount of knowledge about appropriate use of the modal vocabulary into the construction. Either you construct the thing heavily relying on your knowledge of what could, might, must be the case etc.; or it lacks the “right properties” and will not serve “just as well” as the real structure of possibilities, leaving you with wrong predictions for sentences such as “a python can grow longer than eleven meters”.

At best, possible world semantics can be said to clarify the meaning of the modal vocabulary by way of paraphrasing its abundance using one and the same locution in each case, “possible world”. I do not claim that this cannot generate any real understanding. But surely it prompts a question: “So, is that what you basically do in formal semantics? Do you rephrase expressions with synonyms?”

A similar argument suggests itself for other members of the intensional family, such as the tense logic discussed in Gamut (1991, p. 32 ff.). Also here, what seems to be going on is a straightforward rewriting of the time talk into the underlying ontology (in the form of a set of moments and an *earlier than* relation) and “explication” of the former in terms of the latter. “Mary will sing” is true iff there is a moment t' such that the present moment t is earlier than t' and Mary is singing in t' . Again, the only access to these entities, *moments*, is via appropriate time talk. Therefore, in my opinion, the expression on the right is nothing but a clumsy version of saying that Mary will sing at some point, and it is not clear what insight this analysis (or rather paraphrase) provides with respect to the semantics of English tenses.

At this point, one thing is worth making clear: if you consistently occupy yourself with a phenomenon (say, the logical behavior of counterfactual conditionals), with time you come to understand it better, regardless of what formal framework you work with (if any). In that sense, there is no way we can dismiss the whole body of work in formal semantics as trivial or irrelevant to our theoretical knowledge of language. It is however highly questionable whether the real understanding gathered in that project follows from the application of the formal method, or whether it is rather a result of fairly traditional attention to the meaning of expressions, with paraphrase now as ever being the most natural way of semantic clarification. And if we accept that model-theoretic semantics is a project of semantic clarification through paraphrase, one more question comes to mind. Is there an actual confusion that we clarify, and people to whom we

clarify?²⁵ Or do we “clarify” independently of that – just in case?

6.2 Degrees

The problem of ontologizing semantic distinctions, followed by straightforward redescription, is not limited to the intensional models, but seems to permeate through much of the contemporary formal semantics.

In the semantic accounts of gradable adjectives, following Kennedy and McNally (2005) and Rotstein and Winter (2004), the ontology underlying the semantic analysis is enriched with scale-ordered degrees to which an object can have a particular property. A predicate such as *long*, *clean* or *dirty* is then taken to denote, not a function from objects to truth values, but a function from objects to such degrees. In order to account for differences in the logical behavior of gradable adjectives, it is then hypothesized, e.g., that “x is long” is true iff *long* assigns to *x* a *relatively high* degree on the open scale of length; “x is clean” is true iff *clean* assigns to *x* the *maximal* degree on the upper closed scale of cleanness; and “x is dirty” is true if *dirty* assigns to *x* any *non-maximal* degree on the same scale. (It is generally acknowledged that this is a semantic analysis, which needs to be supplemented with an appropriate pragmatics in order to give intuitive results; it is not necessary to dwell on this aspect here.)

Again, I want to argue that this analysis is correct exactly to the same extent to which it is trivial. For saying that semantically, *long* works on an open scale but *clean* expresses the endpoint degree of an upper closed scale is really nothing but a technical paraphrase of the finding that whatever is *long* could also conceivably be *longer*, but an object that is justly called *clean* cannot be made *cleaner* by any more cleaning. (Again, pragmatics aside.) Similarly to the case of possible worlds, the reward of paraphrasing the expressions containing gradable adjectives, comparatives, measure phrases etc. in terms of degrees might be a bit of extra perspicuity, but hardly more than that. One might object here that the very means employed in my alternative paraphrase (such as comparatives or hypothetical clauses) are in need of semantic clarification, and that the basicness of the means of paraphrasing is in the eye of the beholder. But what is perhaps not in the eye of the beholder is that some locutions were (in the very same sense) a part of the ordinary language even before the rise of formal semantics, and some were not.

6.3 Events

The principle of enriching the ontology, as if for free, in the interest of semantic analysis, is particularly well illustrated by Davidson’s (1980a) suggestion of the semantics of events. Here, events are straightforwardly introduced into the ontology, primarily in order to account for the behavior of adverbial modifiers. Basically, “John buttered the toast in the bathroom” is true iff there is a unique

²⁵Has any ordinary semantic confusion ever been resolved by appeal to a formal semantic treatise, in the way in which we consult dictionaries or defer to experts?

toast and a unique bathroom and an event such that it is an event of John's buttering the toast and takes place in the bathroom.

Davidson, unlike many other semanticists, is to be appreciated for providing all sorts of other philosophical reasons for his particular proposal to broaden the assumed ontology (cf. also Davidson, 1980b,c). Also, I do not wish to claim in any way that the standard assumption of a domain of individuals is transparent and unproblematic while anything beyond that (events, degrees etc.) is obscure and undesirable; for Bach (1989, p. 71) is right that “[t]he notion of a thing seems pretty mysterious also”. Still, each step in the way of rewriting the distinctions of everyday talk into the ontology underlying our semantic analysis, by the same token, threatens to be a step towards triviality of that analysis. If enriching the ontology can be a fruitful way in semantics at all, it must be subject to some convincing criteria, and not licensed by the fact that the semantic analysis of some locutions suddenly becomes a piece of cake.²⁶ Such criteria need to follow from a fundamental philosophical discussion.

In the neo-Davidsonian accounts (cf. Maienborn, 2011), the ontological category of events is further subclassified into states, activities, accomplishments and achievements (following Vendler, 1967), in order to account for semantic facts such as that one can *swim for an hour*, but not *swim a mile for an hour* or *explode for a second*. Again, there may well be some useful understanding in this analysis; but to say the least, one wonders where exactly would be the line dividing such useful analyses from purely trivial ones.

In this paper, I hope to have pointed out in somewhat original fashion a handful of problems that, in my opinion, constitute a commitment for anyone engaged in the model-theoretic investigation of natural language semantics nowadays. Formal semantics is *not* a young science any more, and deserves no more special intellectual protection, no more indulgent “we do not yet understand but...”. Once your science grows up, it is time to show how much understanding there actually is under the veil of your technical language...

References

- E. Bach. *Informal lectures on formal semantics*. State University of New York Press, 1989.
- M. Baroni, R. Bernardi, and R. Zamparelli. Frege in space: A program for compositional distributional semantics. *Linguistic Issues in Language Technologies*, 9(6):5–110, 2014.

²⁶Suppose we are to semantically analyse the sentence “Rimbaud's lyrical sensuality is richer than Verlaine's”. Admitting lyrical sensualities as the objects of our ontology, perhaps with the order of richness, would be a fairly absurd way of going about it, I believe.

- R. Brandom. *Making it explicit: Reasoning, representing and discursive commitment*. Harvard University Press, USA, 1998.
- R. Carnap. Meaning postulates. *Philosophical studies*, 3(5):65–73, 1952.
- G. Chierchia and S. McConnell-Ginet. *Meaning and grammar: an introduction to semantics*. MIT Press, 2000.
- C. Chihara. Davidson’s extensional theory of meaning. *Philosophical Studies*, 28(1):1–15, 1975.
- L. Clapp. Davidson’s program and interpreted logical forms. *Linguistics and Philosophy*, 25:261–297, 2002.
- M. Cresswell. *Logics and languages*. Methuen, 1973.
- D. Davidson. The logical form of action sentences. In *D. Davidson: Essays on actions and events*, pages 105–121. Oxford University Press, New York, 1980a.
- D. Davidson. The individuation of actions. In *D. Davidson: Essays on actions and events*, pages 163–180. Oxford University Press, New York, 1980b.
- D. Davidson. Criticism, comment, and defence. In *D. Davidson: Essays on actions and events*, pages 122–148. Oxford University Press, New York, 1980c.
- K. Erk. What do you know about an alligator when you know the company it keeps? Unpublished draft, 2014.
- L. T. F. Gamut. *Logic, Language, and Meaning: Intensional logic and logical grammar*. University of Chicago Press, 1991.
- M. Glanzberg. Semantics and truth relative to a world. *Synthese*, 166:281–307, 2009.
- J. Groenendijk and M. Stokhof. Why compositionality? In G. Carlson and J. Pelletier, editors, *Reference and Quantification: The Partee Effect*, pages 83–106. CSLI Press, Stanford, 2005.
- I. Heim and A. Kratzer. *Semantics in generative grammar*. Blackwell, 1998.
- H. Kamp and M. Stokhof. Information in natural language. In P. Adriaans and J. van Benthem, editors, *Handbook of the Philosophy of Information*, pages 49–112. Elsevier, Amsterdam, 2008.
- C. Kennedy and L. McNally. Scale structure, degree modification, and the semantics of gradable predicates. *Language*, 81:345–381, 2005.
- S. A. Kripke. *Naming and necessity*. Springer, 1972.
- E. Lepore. What model-theoretic semantics cannot do. *Synthese*, 54:167–187, 1983.

- D. Lewis. *Counterfactuals*. Blackwell Publishers, Malden, Massachusetts, 2001.
- N. Maddirala. Philosophy of logical practice: a case study in formal semantics, 2014. Master thesis, ILLC, University of Amsterdam.
- C. Maienborn. Event semantics. In P. Portner C. Maienborn, K. von Heusinger, editor, *Semantics: An international handbook of natural language meaning*, pages 802–829. Walter de Gruyter, Berlin, Boston, 2011.
- D. Marconi. *Lexical competence*. MIT Press, Cambridge, Massachusetts, 1997.
- R. Montague. *English as a formal language*. Ed. di Comunità, 1970.
- B. H. Partee. The development of formal semantics in linguistic theory. In S. Lappin, editor, *The Handbook of Contemporary Semantic Theory*, pages 11–38. Blackwell Reference, 1996.
- J. Peregrin. *Doing worlds with words: formal semantics without formal metaphysics*. Springer, Dordrecht, 1995.
- J. Peregrin. *Meaning and structure: structuralism of (post)analytic philosophers*. Ashgate, Aldeshot, 2001.
- C. Rotstein and Y. Winter. Total adjectives vs. partial adjectives: Scale structure and higher-order modifiers. *Natural Language Semantics*, 12(3):259–288, 2004.
- R. C. Stalnaker. Possible worlds. *Noûs*, pages 65–75, 1976.
- M. Stokhof. Meaning, interpretation and semantics. In D. Barker-Plummer, D. Beaver, J. van Benthem, and P. Scotto di Luzio, editors, *Words, Proofs, and Diagrams*, pages 217–240. CSLI Press, Stanford, 2002.
- M. Stokhof. Hand or hammer? On formal and natural languages in semantics. *The Journal of Indian Philosophy*, 35:597–626, 2007.
- M. Stokhof. The architecture of meaning: Wittgenstein’s Tractatus and formal semantics. In D. Levy and E. Zamuner, editors, *Wittgenstein’s enduring arguments*, pages 211–244. Routledge, London, 2008.
- M. Stokhof. Intuitions and competence in formal semantics. In B. Partee, M. Glanzberg, and J. Skilters, editors, *The Baltic International Yearbook of Cognition, Logic and Communication. Volume 6: Formal Semantics and Pragmatics. Discourse, Context and Models.*, page 1–23. University of Latvia, Riga, 2011.
- M. Stokhof. Formal semantics and Wittgenstein: an alternative? *The Monist*, 96:205–231, 2013.

- J. van Benthem. Suave radicals. In M. Aloni, M. Franke, and F. Roelofsen, editors, *The dynamic, inquisitive, and visionary life of ϕ , $?\phi$, and $\diamond\phi$: A festschrift for Jeroen Groenendijk, Martin Stokhof, and Frank Veltman*, pages 18–21. ILLC, Amsterdam, 2013.
- Z. Vendler. *Linguistics in philosophy*. Cornell University Press, Ithaca, NY, 1967.
- L. Wittgenstein. *Philosophical Investigations*. Blackwell, Oxford, 1967.
- T. E. Zimmermann. Meaning postulates and the model-theoretic approach to natural language semantics. *Linguistics and Philosophy*, 22:529–561, 1999.