Logical Consequence and Natural Language*

Michael Glanzberg
Northwestern University

One of the great successes of the past fifty or so years of the study of language has been the application of formal methods. This has yielded a flood of results in many areas, both of linguistics and philosophy, and has spawned fruitful research programs with names like 'formal semantics' or 'formal syntax' or 'formal pragmatics'.[1] 'Formal' here often means the tools and methods of formal logic are used (though other areas of mathematics have played important roles as well). The success of applying logical methods to natural language has led some to see the connection between the two as extremely close. To put the idea somewhat roughly, logic studies various languages, and the only special feature of the study of natural language is its focus on the languages humans happen to speak.

This idea, I shall argue, is too much of a good thing. To make my point, I shall focus on consequence relations. Though they hardly constitute the full range of issues, tools, or techniques studied in logic, a consequence relation is the core feature of a logic. Thus, seeing how consequence relations relate to natural language is a good way to measure how closely related logic and natural language are. I shall argue here that what we find in natural language is not really logical consequence. In particular, I shall argue that studying the semantics of a natural language is not to study a genuinely logical consequence relation. There is indeed a lot we can glean

*Thanks to Jc Beall for valuable discussions of the topics explored in this paper. Versions of this material were presented at the Conference on the Foundations of Logical Consequence sponsored by Arché, St. Andrews, 2010, and at Tel Aviv University, 2011. Thanks to all the participants there for many helpful questions and comments, and especially to Colin Caret, Josh Dever, Christopher Gauker, Ole Hjortland, Graham Priest, Agustín Rayo, Stephen Read, Stewart Shapiro, and Brian Weatherson. I am especially grateful to my commentator at St. Andrews Derek Ball, and to an anonymous referee, for incisive comments.

[1]For instance, Portner and Partee (2002), Sag *et al.* (2003), and Kadmon (2001).

about logic from looking at our languages, and at our inferential practices. But, I shall argue here, we only get to logic *proper* by a significant process of *identification*, *abstraction*, and *idealization*. We first have to identify what in a language we will count as logical constants. After we do, we still need to abstract away from the meanings of non-logical expressions, and idealize way from a great many features of languages to isolate a consequence relation. This process takes us well beyond what we find in a natural language and its semantics. We can study logic by thinking about natural language, but this sort of process shows that we will need some substantial extra-linguistic guidance—some substantial idea of what we think logic is supposed to be—to do so. We do not get logic from natural language all by itself.

This is not a skeptical thesis, about logic or about language. It accepts that there are substantial facts about what makes a relation a logical consequence relation, and it accepts there are substantial facts about the semantics of natural language. Nor is it a brief against any particular methods in the study of language. Logical methods, as I mentioned, have proved their worth in studying language many times over. It is, rather, an autonomy thesis: the two sets of facts are fundamentally autonomous, thought the processes of identification, abstraction, and idealization can forge some connections between them.

There is one large proviso to the conclusions I just advertised. Defending them in their strongest form will require assuming a fairly restrictive view of what logical consequence relations can be like, that distinguishes them from other related notions. This assumption is entirely consonant with a long tradition in logic. It is certainly the way logic was thought about in the work of Frege and Russell, and others at the turn of the twentieth century when the foundations of modern logic were being laid. But in spite of its lofty pedigree, this view is not universally shared, and indeed there is an equally vaunted tradition in logic that rejects it. So, we should not take such a view for granted. Even so, my goal here is not to defend any particular view of the nature of logic, but rather to see how natural language relates to logic given our views on logic's nature. So, I shall begin by assuming a highly restrictive view of the nature of logical consequence, and defend my conclusions about natural language and logical consequence accordingly. I shall then reconsider my conclusions, in light of a more permissive view. I shall argue that the conclusions still hold in the main part, but that extremely permissive views might have some ways to avoid them. I shall suggest, however, that such extreme views run the risk of buying a close connection between logic and natural language at the cost of making it uninteresting, or even trivial.

My discussion in this paper will proceed in five sections. In section I, I shall motivate the idea that logic and language are closely connected, and spell out my contrary thesis that they are not. I shall also briefly discuss some ideas about logic and related notions I shall rely on throughout the paper. In section II, I shall offer my main argument that logic and natural language are not so closely connected. In doing so, I shall articulate a little of what I think the semantics of a natural language is like, and show how entailments, but not logical consequence relations, are found in natural language. The arguments of section II will presuppose a restrictive view of logical consequence. In section III, I shall reconsider those arguments from a permissive view of logical consequence. I shall argue that my main conclusions still hold, though perhaps in weakened form, and some important qualifications are in order. I shall then show in section IV how one can move from natural language to logical consequence by the three-fold process of identification, abstraction, and idealization. Finally, I shall offer some concluding remarks in section V.

I Preliminaries and Refinements

It is very common, at least in some circles, to speak of a logic and a language in the same breath. This perhaps makes the idea that logical consequence and semantics of natural language are closely related an inviting one. The principal thesis of this paper is that they are not so closely related. However, this thesis stands in need of some refinement, and some important qualifications.

To provide them, I shall begin by briefly articulating the perspective that sees logic and language as closely connected. I shall then discuss the notion of logical consequence itself, and distinguish more permissive and more restrictive views on the nature of logical consequence. That will allow us to distinguish (restrictive) logical consequence from some related notions. With these preliminaries in hand, I shall be able to formulate a more refined version of my main thesis. I shall conclude this section with some discussion of what the thesis implies about the application of formal methods to natural language semantics.

I.1 Logics and Languages

For many logicians, it is languages, i.e. formal languages of particular sorts, that are the primary objects of study. A fairly typical example is the textbook of Beall and van Fraassen (2003), which studies formal languages comprised of a formal syntax, a space of valuations of sentences of the syntax, and a relation of satisfaction between sentences and valuations. Logical consequence is preservation of satisfaction (in most cases, preservation of designated value). This particular textbook focuses on sentential logic, so the valuations are typically determined by assignments of truth values to atomic sentences, but they could very well be determined by models of the right sorts, and for quantificational languages, they will be.

Making such languages the basic elements of logic is especially convenient for the study of a variety of logics, as it gives a natural unit of study that can vary. We ask about different languages, and explore and compare their logical properties. But the connection might well go deeper, and in some ways, it must. Genuine logical relations are interesting not just for their abstract properties, but what they tell us about connections between meaningful sentences, which express what we think. Conversely, the basis for a consequence relation is often thought to be found in the meanings of sentences. For instance, Beall and van Fraassen (2003, p. 3) write:

> Logic pertains to language. In a logical system we attempt to catalog the valid arguments and distinguish them from the ones that are invalid. Arguments are products of reasoning expressed in language. But there are many languages and different languages may have a different structure, which is then reflected in the appropriate logical principles.

When we look for valid arguments, we look at meaningful language as the medium of expressing them. But moreover, we think of the meanings of sentences as crucial for determining whether or not an argument is valid. Thus, we tend to think of the languages that provide consequence relations as genuine languages, capable of expressing thoughts in much the ways our natural languages do.

Combining these ideas, we can formulate two theses. The first is the *logics in formal languages thesis*:

> Logical consequence relations are determined by formal languages, with syntactic and semantic structures appropriate to

isolate those relations.

The logics in formal languages thesis is not entirely trivial, but it is not particularly controversial either. Working in terms of formal languages is one theoretical choice among many on how to develop important logical notions. Like any theoretical choice, it has various consequences that might be weighed differently, but there is little room to say it could be outright wrong. Perhaps more controversial is the idea that there are multiple logics associated with multiple languages. This might be challenged, and strong ideas about logical pluralism are indeed controversial. I shall return to some issues surrounding this point in a moment. But for now, I shall simply observe that the logics in formal languages thesis is optional, but by itself not very contentious. I formulate it to help frame the next thesis, which will be genuinely contentious.

The key idea of the logics in formal languages thesis is that formal languages are the unit of study for logic, and so, formal languages must determine consequence relations. A formal language is a bundle of elements, usually containing something like a synax, a space of valuations or models, and a relation of satisfaction. A consequence relation is definable from these, typically as preservation of designated value. Variation in ways of presenting formal languages will allow for variation in how consequence relations are defined on them, but the important idea is that formal languages contain enough elements that consequence relations can be defined on them alone, just as we see in the typical case. Thus, the logics in formal languages theses holds that consequence relations are *in* formal languages, in the sense that they are definable from them. I shall likewise sometimes talk about languages *containing* consequence relations.[2]

As we just discussed, it is inviting to think that these formal languages are importantly like natural language. How alike? I suggest the key idea is that a natural language shares important logical features with formal languages. Most importantly, they share the feature of containing logical consequence relations. Thus, studying a range of formal languages may expand our horizons, but it shows us more of what we already can find in natural language. This leads to our second thesis, the *logic in natural language thesis*:

[2]Thus, the claim that a language contains a logic is much stronger than the claim that we can talk about a logic in the language. We can talk about all sorts of things in languages that are not definable on the language itself. We can talk about physics in a language too, but I doubt any formal or natural language contains physics.

> A natural language, as a structure with a syntax and a semantics, thereby determines a logical consequence relation.

The syntax of natural languages differ from that of our favorite formal languages, and in some ways, perhaps, their semantics does too. But regardless, according to the logic in natural language thesis, they are just more languages, and determine logics just like formal languages do.[3]

Whereas the logics in formal languages thesis is relatively banal, the logic in natural language thesis is one I shall argue against. But, just how strong the thesis is, and what it takes to argue against it, depends on how we view consequence relations. Before we can get an accurate statement of the thesis, and an adequate appreciation of what rejecting it amounts to, we must consider what counts as a logical consequence relation.

I.2 Logical Consequence

In this section, I shall briefly review some ideas about logical consequence. My goal here is limited. I shall not try to defend a view of the nature of logical consequence; rather, I shall try to survey enough options to better frame the logic in natural language thesis.

At its core, logic is the study of valid arguments, as Beall and van Fraassen say above. Of course, this is not all that logicians study, as the many topics in such areas as set theory, model theory, or recursion theory make clear. But this is the core feature that makes logic logic. Logical consequence is that relation that makes arguments valid: it is the relation that holds between a set of sentences and a sentence when the first set comprises the premises and the second sentence the conclusion of a valid argument.[4] But then the important question about the nature of logical consequence is simply what makes an argument valid? There are a few key ideas that I shall suppose for argument's sake. As I mentioned, my aim here is not to provide a definitive

[3] This attitude is perhaps most strongly expressed by Montague (1970, p. 222), who writes "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." At least, so long as we take the semantics of artificial languages to include a consequence relation, then Montague's view includes, and goes well beyond, the logic in natural language thesis.

[4] Of course, things can be complicated here in various ways, but we will not reach a level of detail where such complications would matter in this discussion.

analysis of the notion of logical consequence—that would be far too hard a task. Rather, I simply want to say enough to distinguish logical consequence from some of its neighbors.

Perhaps the main feature of logical consequence is what is sometimes called *necessity*: if S is a consequence of a set X of sentences, then the truth of the members of X necessitates the truth of S, or equivalently, it is impossible that each element of X be true and S be false. What makes an argument valid, necessity notes, is in part that the conclusion cannot be false while the premises are true. This is imprecise, as the notion of necessity at work is not spelled out. Even so, it is sufficient to distinguish logical consequence from, for instance, inductive support. Even a sentence which enjoys strong inductive support on the basis of some assumptions can fail to be true while the assumptions are true.[5]

Necessity is one of the main features of any relation we would call a logical consequence relation. Another, especially in the tradition of Tarski (1936) or Quine (e.g. Quine, 1959, 1986), is what is sometimes called *formality*. Formality is hard to spell out in full generality, but the main idea is that logical consequence is somehow a 'formal' relation, holding in virtue of the forms of the sentences in question. This condition is typically taken to rule out implications like 'John is Bill's mother's brother's son, therefore, John is Bill's cousin' as not genuinely logical. It is not, the reasoning goes, because it relies for its validity on the specific meaning of 'cousin' rather than the formal properties of the sentence. Of course, how formality is implemented depends on just what we take the relevant 'formal' structure of a sentence to be. Whereas necessity is likely to be recognized as a requirement on consequence relations across the board, formality is more contentious, and how far it will be accepted will depend on how it is spelled out.

One leading approach to spelling out formality, since Tarski if not earlier, has been to identify special logical structure, and propose that logical consequence must hold in virtue of only that structure. It is thus formal in that it holds in virtue of specified structure or form in sentences. In the post-Tarskian tradition, we often capture this by insisting that there are privileged *logical constants* in sentences, and logical consequence holds in virtue of their properties. Spelled out in a model-theoretic way, the idea is that logical consequence holds solely in virtue of the meanings of the logical constants, and

[5] My discussion of logical consequence relies heavily on those of Beall and Restall (2009) and MacFarlane (2009).

hence, we hold those meanings fixed, but allow all other meanings to vary, as we work out model-theoretic consequence relations.

Necessity and formality may well interact. For instance, they do in the standard post-Tarskian model-theoretic view of consequence. This view relies on a range of models to characterize the consequence relation. The models thus provide the possibilities which give substance to necessity. They also implement formality, by allowing the meanings of all the non-logical terms of a language to vary freely, and thereby single out logical constants whose meanings underlie validity. They thus offer a combined version of both formality and necessity. The effect of this combination is enough to rule out familiar analytic entailments as not logical consequences. The 'cousin' inference is an example, and indeed, so is 'Max is a bachelor, therefore, Max is unmarried', which also fails to be a logical consequence on many views.[6]

When implemented this way, necessity and formality are often understood as conspiring to make logical consequence a very narrow notion. We see this, for instance, in the way they rule out analytic entailments as not logical. Behind this conclusion is a very general attitude towards logic, which holds there are substantial constrains on what makes for genuine logical consequence, and finds that only such a narrow notion meets those constraints. This attitude is a long-standing one in the philosophy of logic. It is most prominent in the strand of logic starting with late 19th century developments in the foundations of mathematics in work of Frege and Russell (e.g. Frege, 1879; Russell, 1903; Whitehead and Russell, 1925–27), going through such figures as Gödel (1930) and Skolem (1922), and then moving on to Tarski (1936) and Quine (1986), to more recent work by Etchemendy (1990), Shapiro (1991), and Sher (1991), among many others. This tradition is not uniform in its views of what grounds logical consequence (for instance, Frege and Quine and Tarski would see the formality constraint very differently). But it is uniform in thinking that there is some important underlying notion of logic, and that logic plays particular roles and has a particular status. Its special epistemological status was especially important in discussions of logicism, and its metaphysical status in discussions of ontological commitment.

[6]There is considerable historical debate over just what Tarski's view of logical consequence was, sparked, in large part, by Etchemendy (1988). For this reason, I have talked about the 'post-Tarskian view', which is embodied in standard contemporary model theory. This view no doubt stems from work of Tarski, whether it is Tarski's original view or not. For a review of some of the historical work on this and related issues, see Mancosu (2010).

Both features are apparent in the discussions of the role of logic in the foundations of mathematics as first order logic emerged in the early 20th century. This tradition sees substantial constraints on what makes a relation logical consequence, which flow from the underlying nature of logic. These constraints substantially restrict what can count as logic. This tradition is thus, as I shall call it, *restrictive*, in that it provides highly restrictive constraints answering to a substantial underlying notion.

There is another, more *permissive*, tradition in logic, going hand-in-hand with work on 'non-classical' or 'non-standard' logic. This tradition also has a vaunted pedigree. The range of ideas and issues that it encompasses is very wide, but some of them can be traced back to Aristotle, and were important in the lively medieval logical literature.[7] The permissive tradition does not necessarily abandon the idea that there is some underlying notion of logical consequence, but it interprets the idea much more broadly, and sees much more variety in how the constraints on logical consequence might be applied. The result is a willingness to consider a range of logics as at least candidates for showing us logical consequence relations. It is perhaps natural to think of the permissive approach as going hand-in-hand with *logical pluralism* (e.g. Beall and Restall, 2006). If it does, then it will think that more than one logic genuinely is logic. But permissive views certainly need not be completely indiscriminate. For instance, many logicians in the relevance tradition have doubts about whether classical logic really is logic, often driven by the so-called paradoxes of implication.

Permissive views need not reject the ideas of necessity and formality, and notably, Beall and Restall's logical pluralism does not. But they will interpret these notions more expansively than many classically oriented restrictive views do. One way they can do so is to expand or contract the range of possibilities which inform the necessity constraint, or modify the way truth in a circumstance is characterized. They can also vary the range of logical constants which give content to formality. Both have been done, many times for many different reasons. Often some combination of both proves fruitful.

Another axis on which we might compare views of logical consequence is according to how broad or narrow a notion of consequence they accept. Classical logic, for instance, is narrower than a consequence relation that

[7] I am not qualified to review this history, so I shall defer to experts such as Kneale and Kneale (1962), Kretzmann *et al.* (1982), Read (in press), Smith (1995), and the many references they cite. For a review of more contemporary developments, see Priest (2008) and the many references therein.

includes analytic entailments like the 'bachelor' entailment. These issues are substantially independent of those of permissive versus restrictive views. How broad or narrow the notion(s) of consequence you accept are is determined by just which constraints you impose, not the general issue of permissive versus restrictive approaches. Even so, we may expect permissive views to more readily entertain a range of broader consequence relations than restrictive views do. Likewise, in introducing the restrictive view, I noted that it tends to indicate narrow notions of logical consequence. It will thus simply our discussion to assume that restrictive views are committed to only a narrow notion of consequence, while permissive views can entertain broad ones as well. This reflects a trend in thinking about consequence, but it is a simplification. For one reason, it leaves out that permissive views might well entertain more narrow consequence relations than restrictive views do. But, it will be a useful simplification to make.

For discussion purposes, I shall usually assume that the restrictive view is going to opt for something like classical logical consequence, as in fact the tradition I associated with the restrictive view did opt for; while permissive views can consider consequence relations that reflect other sorts of entailments. I shall not worry here about differences within classical consequence relations, like first versus second order logics.[8] In light of the logics in formal languages thesis, and the way we have glossed the constraints of necessity and formality, it will be natural to assume a model-theoretic account of consequence relations, both classical and otherwise. This assumption will set up the most likely route to the logic in natural language thesis, so it is a harmless one to make here.[9]

I.3 Implications and Entailments

We now have at least roughly sketched some ideas about logical consequence, and distinguished permissive from restrictive views of consequence. In what

[8] Type theories are used in a great deal of work in semantics, but as we will see in section II, not in a way that directly indicates a consequence relation.

[9] I thus have relatively little to say about proof-theoretic accounts of logical consequence, for instance, as explored in work of Dummett (e.g. Dummett, 1991) and Prawitz (e.g. Prawitz, 1974), all building on seminar work of Gentzen (1935). For a recent survey, see Prawitz (2005). In that survey, Prawitz explicitly endorses the general constraints of necessity and formality, though of course, not the model-theoretic gloss on them I typically employ here.

follows, I shall argue against the logic in natural language thesis on the basis of those ideas, especially on the basis of restrictive views of consequence. But when it comes to natural language, it will be important to distinguish logical consequence from potentially broader related notions, such as we have already seen with analytic entailments. In this section, I shall review some of these notions.

To fix terminology, let us start with *implication*. I shall take this term to be very broad, covering many relations between sentences including not only logical consequences and subspecies of them, but looser connections like those captured by defeasible inference. Following the philosophy of language tradition, we might also see implications between utterances of sentences, often defeasible, such as the one discussed by Grice (1975) that typically obtains between 'There is a gas station around the corner' and 'It is open'.

By restrictive lights, implication is much broader than logical consequence, but it is such a wide and loose notion that it is not clear if it is apt for formalization even by a very broad notion of consequence. Permissive approaches have done substantial work on some species of it, notably defeasible inference.[10] Some attempts have been made to make rigorous the computations that might support implicatures, but they tend to focus more on computation than consequence per se.[11]

Within the broad category of implications, two specific notions will be important. One, narrow logical consequence (i.e. classical first order consequence or something thereabouts), we have already seen. The other is *entailment*. I shall understand entailment as a truth-conditional connection: P entails Q if the truth conditions of P are a subset of the truth conditions of Q. The usual modification for multiple premises holds. We already have seen enough to know that entailment is a wider notion than narrow logical consequence. Analytic entailments like the 'cousin' implication above are entailments by the current definition, but not narrow logical consequences. By many lights, entailments go beyond analytic entailments. They will for instance, if truth conditions are metaphysically possible worlds. If so, and assuming Kripke-Putnam views of natural kind terms, then 'x is water, therefore x is H_2O' is an entailment. We will encounter more entailments as we proceed. What we need now is simply that entailment is a markedly broader notion than narrow logical consequence, though narrower than some notions

[10]See, for instance, Antonelli (2005). See Horty (2001) for a survey of related ideas.
[11]See, for instance, Asher and Lascarides (2003) and Hirschberg (1985).

of implication, including implicatures like the 'gas station' one.

We now have seen three related ideas. A very broad notion of implication, a very narrow notion of logical consequence, and an intermediate notion of entailment. Permissive views of logic have done extensive work to capture notions of entailment as broad consequence relations. Indeed, such work has identified a number of distinctions within the category of entailment.[12] As I noted above, permissive views have also attempted to capture some aspects of the wide notion of implication. So, we should not think these relations to be beyond the range of permissive approaches to logical consequence, so long as they are broad enough in their notions of consequence. They are, however, clearly beyond the range of the restrictive view, and fail to offer narrow consequence relations. We will see in section II.2 that natural language presents us with a striking variety of entailments, but I shall argue, not narrow logical consequence.

I.4 The Refined Thesis

Now that we have some preliminaries about logic and related notions out of the way, we can return to the main claims of this paper. Above I formulated a thesis of logics in formal languages that claims that consequence relations are determined by formal languages. (Both the permissive and restrictive approaches are compatible with this thesis.) But the important thesis for this paper is the more contentious logic in natural language thesis, which holds that a natural language, construed as including a syntax and a semantics, determines a logical consequence relation. The main contention of this paper is that the logic in natural language thesis is false. I shall argue that it is clearly false if we adopt the restrictive view of logical consequence. There are no doubt entailment relations in natural language, determined by the semantics of a language, and there are many other implication relations as well. But the semantics and syntax of a natural language does not determine what the restrictive view of consequence takes to be a logical consequence relation.

If we adopt the permissive view, this claim becomes rather more nuanced. I shall argue that we still have good reason to think the logic in natural language thesis is false, even if we adopt a permissive view. However, I shall also grant that there are some, perhaps extreme, permissive views that might

[12] See, for instance, Anderson and Belnap (1975) and Anderson *et al.* (1992).

support some forms of the thesis. Even so, I shall argue, they run the risk of stretching the notion of logic too far, and thereby undercutting the interest and importance of the thesis.

Though I shall argue against the logic in natural language thesis here, I shall not claim there is no connection between logic and natural language. We can glean some insight into logical consequence, and indeed even narrow logical consequence, by studying natural language. The reason is that the entailments and other implications we do find in our languages, and our wider inferential practices, provide a rich range of examples around which we can structure our thinking about logical consequence. But to do so correctly, we must get away from the entailments and implications of a human language and human inferential practice, and isolate genuine logical consequence. I shall argue that the strategy of identification, abstraction, and idealization I mentioned above is a useful one for taking this step, given the kinds of information that natural languages really do provide for us. I shall argue that even the permissive view needs to take these steps, to get theoretically substantial logics out of natural language. It may be that for the permissive view, some of the steps may be shorter than those the restrictive view needs to take, but the same kinds of steps must be taken by both.

I.5 No Logic in Semantics?

I am proposing that in a narrow sense, natural language has no logic, and I thereby echo Strawson's famous quip (Strawson, 1950). Though I echo some of the letter of Strawson, I do not follow the spirit. To make this vivid, let me spell out several things I am not claiming. First and foremost, if we adopt a restrictive view and accept only narrow consequence as consequence, saying we do not find logical consequence in natural language does not by any means say that natural language is immune to study by formal or mathematical methods.[13]

In fact, we will see a number of places where we find that natural language semantics makes productive use of tools and techniques from logic. I

[13] The original from Strawson (1950, p. 27) reads, "Neither Aristotelian nor Russellian rules give the exact logic of any expression of ordinary language; for ordinary language has no exact logic." I am not sure exactly what Strawson has in mind here, but it is common to read him as advocating the view I reject, that natural language is immune to study by formal methods. As I like many of the arguments of Strawson's paper, I often tell my student to pay close attention to every part of it except the last line.

shall explain as we proceed how this can happen without indicating a narrow consequence relation in natural language, and indeed, seeing how this happens will provide good reasons for rejecting the logic in natural language thesis. Once we see how we really use logic in the study of natural language semantics, the thesis loses its intuitive appeal.

The applicability of logical methods to the study of language, in spite of the failure of the logic in natural language thesis, should not itself be surprising. In singling out the notion of logical consequence as the core of logic, we should not be blind to the impressively wide range of applications of logical (and more generally, formal) methods, which goes well beyond the study of logical consequence per se. Logical structure, in the rough sense of what is tractable via the methods of logic, can be found in many places. Computer science finds it in the organization of data and computation, linguistics finds it not only in some aspects of semantics, but in the syntactic organization of parts of sentences. Sociologists find it in the organization of social networks. Indeed, some of the core underlying structures of logic, like Boolean algebras, seem to be found practically anywhere you look for them.

From a restrictive point of view, many of these applications go beyond the study of logical consequence. For instance, the syntactic structure of human language seems clearly not to be a matter of logical consequence according to a restrictive approach, even if we can represent substantial portions of it with the formalisms of logic via the Lambek calculus (Lambek, 1958; van Benthem, 1991; Moortgat, 1997). From a very permissive point of view, say, one which is happy to talk about a logic of syntax via the Lambek calculus, things may look somewhat different. But as I said above, I shall argue they are not all that different. We do not get to clearly articulated broad consequence relations in studying language without departing from the on-the-ground study of language, even if our methods are formal ones.

Denying the logic in natural language thesis in no way argues we should put aside logic when we come to study natural language. Rather, it argues that we should see logical consequence proper and the semantics of natural language as substantially autonomous, but linked by such processes as abstraction and idealization.

II The Logic in Natural Language Thesis from the Restrictive Point of View

We now have a more careful articulation of my main claims, that the logic in natural language thesis fails, but that natural language and logic can be connected by the three-fold process of identification, abstraction, and idealization. My defense of these claims will come in three parts. First, in this section, I shall argue against the logic in natural language thesis assuming a restrictive view of logical consequence. In the next section III, I shall reconsider those arguments from a permissive point of view. I shall then turn to how to bridge the gap between natural language and logic in section IV.

My discussion of the logic in natural language thesis here will present three arguments. The first will argue against the thesis directly, by showing that the semantics of natural language does not provide a consequence relation. The second and third will show that there is no way around this conclusion. The second argument will show that the implications natural language does provide to us are not generally logical consequences. The third will show that natural language does not distinguish logical constants, and so formality cannot be read off the structure of natural language. Thus, we find no consequence relation in the semantics of natural language proper, and cannot find one encoded in natural language by more indirect means. Throughout this section, I shall assume a restrictive view of consequence without further comment.[14]

II.1 Semantics, Model Theory, and Consequence Relations

In this section, I shall argue that consequence relations are not provided by the semantics of natural language the way they are provided by formal languages (assuming the logics in formal languages thesis). This is so, I shall argue, even assuming a truth-conditional approach to semantics, and even assuming its model-theoretic variant. The model theory we do for semantics

[14]To come clean, I am inclined to take a restrictive view, and claim that something like classical logical consequence is the right notion (though I am not really decided on issues like those of first versus second order logic, the status of intensional logic, etc.). But it is not my goal to defend any such views here.

is not the sort of model theory that provides model-theoretic consequence relations (and might be better called something other than 'model theory'). To argue this, I shall first lay out a little of what I think the semantics of a natural language is like. This will show that a viable natural language semantics, in particular, a viable truth-conditional semantics, cannot provide a model-theoretic consequence relation. I shall then explain why this is the case in spite of the apparent use of model theory in so-called 'model-theoretic semantics', and discuss some aspects of how model-theoretic techniques can shed light on natural language without providing consequence relations.

II.1.1 Absolute and Relative Semantics

Many logicians, lured by the siren song of natural language, have found themselves thinking of model theory as applying to fragments of natural language. This trend got a huge boost in the 1960s and 1970s with developments in the model theory of intensional logics, which made applications to natural language easy to find (e.g. Montague, 1968; Scott, 1970). Eventually, Montague (e.g. Montague, 1970, 1973) in effect proposed that to do the semantics of natural language includes doing its model theory (and proposed that it could really be done). More might be required, e.g. Montagovian meaning postulates. But along the way to giving a semantics in the Montagovian mold, you will provide a whole space of models, and determine truth in those models for sentences, and so, you will have done your model theory. If this was right, then the logic in natural language thesis would be sustained (though we would want to check that the result lived up to restrictive standards). Indeed, this is the main reason the logic in natural language thesis might be thought to be correct.[15]

I shall argue that this is a mistake. This will show that the Montagovian route to the logic in natural language thesis fails. Moreover, it will show

[15] For more recent presentations of Montague Grammar, see for instance Dowty *et al.* (1981) or Gamut (1991). Other important early contributions along similar lines include Cresswell (1973) and Lewis (1970). Of course, these works did not appear in a vacuum, and earlier works of Frege (e.g. Frege, 1891), Tarski (e.g. Tarski, 1935), Carnap (e.g. Carnap, 1947), Church (e.g. Church, 1940), Kripke (e.g. Kripke, 1963), and others stand as important predecessors.

Though Montague and a number of these other authors link logic and natural language, this marks a departure from the main trend of early work in the restrictive tradition I discussed in section I. Frege and Russell, for instance, saw only distant connections between logic and natural language.

that the semantics of natural language does not build in a consequence relation in anything like the way the Montagovian route supposes. I take this to be a good reason to reject the logic in natural language thesis, though I shall provide two other reasons below. My argument goes by way of a reconsideration of an old debate about model-theoretic semantics for natural language. The debate, between neo-Davidsonian advocates of so-called 'absolute' truth-conditional semantics challenged the idea that model-theoretic techniques could be used to study natural language at all.

The debate takes place within a program of truth-conditional semantics for natural language. The goal of this program is to provide an account of a key aspect of speakers' linguistic competence. In particular, semantics seeks to provide an account of what a speaker understands when they grasp the meanings of their sentences, i.e. what they know when they know what their sentences mean.

I shall take it for granted that truth conditions provide a central aspect of this knowledge. A key part of what a speaker knows when they know the meaning of 'Snow is white' is precisely that it is true if and only if snow is white. This is a non-trivial assumption, rejected, for instance, by a number of conceptual role or cognitive approaches to semantics. I think it is correct, but I shall not defend it here. Rather, I shall take it a starting point for asking about how logic, and particularly consequence relations, might find their way into natural language.[16]

[16] For discussion of this role for truth conditions, see for instance Higginbotham (1986, 1989b), Larson and Segal (1995), and Partee (1979).

The important question, of course, is whether this assumption is correct; and on that point, I do not have much to add to the current literature, and this is certainly not the place to pursue the issue in depth. But, one might wonder if making this assumption makes my whole argument somewhat parochial. I do not think it does. At least in empirically-minded work in semantics and related work in philosophy, truth-conditional semantics is a well-established research program. It might well be the dominant one (I think it is), though it is not my intention to dismiss work in cognitive semantics, which has been important in the literature on lexical semantics. Whether it is the dominant research program or not, truth-conditional semantics certainly enjoys a sufficiently important place in research in semantics to make asking about its connections to logic an important one. If, as I believe, assuming a truth-conditional perspective is correct, then all the more so. In the less empirically-minded literature in philosophy of language, there has been more discussion of conceptual role theories of meaning, inferentialist theories of meaning, and use theories of meaning. But, these have never really gotten off the ground as empirical theories. Of course all these deserve more discussion, but from the empirically-minded perspective I am adopting here, they to not really provide well-developed alternatives to

From this starting point, it is no surprise that logicians have sometimes found a very close connection between semantics and model theory, and through that, a connection with logical consequence, just as the logic in natural language thesis would have it.[17] Models, after all, seem just right to play the role of individual conditions under which the sentences of a language are true. So, to specify the truth conditions of a sentence, we might suppose, is precisely to specify the class of models in which it is true. This is one of the key components of Montague's own approach to the semantics of natural language (e.g. Montague, 1973).[18]

On this view, a good semantic theory will, among things, assign sets of models to sentences, which are taken to capture their truth conditions. Of course, we want more than that; not least of which, we want our theory to correctly derive the truth conditions of sentences from the meanings of their parts (we presumably want some form of *compositionality* to hold). Putting these two together, a reasonable requirement for a theory which will assign sets of models as truth conditions is that it can derive results like:

(1) For any model $\mathfrak{M}$, 'Ernie is happy' is true in $\mathfrak{M} \iff \text{Ernie}^{\mathfrak{M}} \in \text{happy}^{\mathfrak{M}}$.

---

truth-conditional semantics.

[17] Though in many cases, logicians are more interested in the pure mathematics than the psychology. This was certainly Montague's attitude. As Thomason (1974, p. 2) puts it, "Many linguists may not realize at first glance how fundamentally Montague's work differs from current linguistic conceptions. Before turning to syntactic theory, it may therefore be helpful to make a methodological point. According to Montague the syntax, semantics, and pragmatics of natural languages are branches of mathematics, not of psychology." (For further discussion, see again Partee (1979), or Zimmermann (1999).) Perhaps this was not realized by many, and it is perhaps no surprise that authors such as Higginbotham identified with the neo-Davidsonian tradition in semantics. But whether intended by Montague or not, some aspects of his apparatus have been incorporated into an approach to semantics within the broader tradition of generative linguistics, as is witnessed by the textbooks of Chierchia and McConnell-Ginet (2000) and Heim and Kratzer (1998). I shall briefly discuss what of Montague's apparatus got so-incorporated below, but a more full discussion shall have to wait for other work in progress. Finally, I should note that the general idea that semantics has something to do with what we know when we know a language can be completely divorced from the Chomskian perspective of Higginbotham or Larson and Segal, as in work of Dummett (e.g. Dummett, 1991).

[18] In sketching Montague's ideas, I am suppressing a great deal of detail that is not important for the issue at hand. For instance, Montague's approach also relies on other elements, including categorial grammar, intensional type theory, and meaning postulates, but these do not change the basic place of models in the theory.

(Ernie$^{\mathfrak{M}}$ is the extension (or other appropriate value) of 'Ernie' in $\mathfrak{M}$.) If we can do this for a large fragment of a human language, compositionally, the view holds, we have thereby elaborate a good semantic theory.

In many applications, the models involved are modal, and contain a set of possible worlds, as in the classic Montague (1973) or Lewis (1970). But the assignment of semantic values is still done relative to a model, as well as to a world, and sometimes a time. We will not be worried here about intensional constructions, so we can just think about assigning reference and truth in a model. In the long run, the use of intensional model theory might well affect what consequence relations are at issue, but it will not affect the basic route from truth conditions to models to consequence, so we can safely ignore this issue here.

The model-theoretic or Montagovian approach to truth-conditional semantics is one of two classic approaches. The other, following Davidson (1967) (who himself in some ways follows Tarski 1935), emphasizes deriving clauses like:

(2) 'Ernie is happy' is true $\Longleftrightarrow$ Ernie is happy.

Davidson, following Quine (e.g. Quine, 1960), emphasized the extensional nature of such a theory, and also its non-model-theoretic pedigree. No model is mentioned, and clauses like these are typically derived from statements of reference and satisfaction properties, like that 'Ernie' refers to Ernie. But what is most important for us is that the resulting T-sentences or disquotational statements state the truth conditions of sentences.[19]

It may look like these two approaches do essentially the same thing. One is more proof-theoretic, emphasizing theories that can derive canonical statements of truth conditions. The other is more model-theoretic, explicitly referring to models. Yet both seem to be in the business of working out how the truth conditions of a sentence are determined by the reference and satisfaction properties of its parts. Both thereby hope, in light of the assumptions we have made about meaning, to represent some important aspects of a speaker's knowledge of meaning, i.e. their semantic competence.

In spite of this, it is often thought that the two approaches to semantics are very different. In fact, it has been argued, notably by Lepore (1983), that model-theoretic semantics is somehow defective, or at least less satisfactory

[19]Davidson himself did not ascribe to some of the assumptions about linguistic competence I made above, but some neo-Davidsonians do, including Higginbotham and Larson and Segal cited above.

than absolute semantics (cf. Higginbotham, 1988). Far from being variants on the same basic idea, Lepore argues, the model-theoretic approach has built-in failings that make it inappropriate for doing semantics at all, and it is hardly equivalent to a neo-Davidsonian semantics.

Though I do not think the morals to be drawn from this argument are what many neo-Davidsonians do, there is something importantly right about the argument, and it will reveal something important about the connections between logical consequence and truth-conditional semantics. This will lay the groundwork for rejecting the logic in natural language thesis.

Lepore's main point is that model-theoretic semantics can only provide *relative* truth conditions: i.e. conditions for truth in or relative to a model. And, you can know those and not know what the sentence means. You can know that for any model, 'Snow is white' is true in that model if the extension of 'white' in that model includes the referent of 'snow', without having any idea what the sentence means. You could have no idea that it talks about snow, or whiteness. It is no better than knowing that 'The mome raths outgrabe' is true in a model if the extension of 'raths outgrabe' in the model includes the referent of 'the mome' in the model. We know that, but (I at least) have no idea what this sentence means.[20]

Davidsonian or *absolute* statements of truth conditions, of the kind we get from techniques stemming from Tarski's work on truth (Tarski, 1935), do tell you much more. They tell you the sentence 'Snow is white' is true if and only if snow is white, which is what we wanted. As you do in fact understand what 'snow' and 'white' mean, this tells you much more than that in any model, some element falls in some extension. It tells you that this stuff—snow–has this color—white. Similarly, we cannot, in a language we understand, write down a T-sentence for the Lewis Carroll 'The Mome raths outgrabe'. Hence, Lepore argues, the model-theoretic approach to semantics fails to characterize enough of what speakers know about the meanings of their words and sentences, while the absolute or neo-Davidsonian approach does much better.

I think the conclusion of this argument is correct. It follows that we do not provide a consequence relation as part of the semantics of natural language. In particular, it follows that we do not provide a model-theoretically defined consequence relation, by providing a space of models and a truth in a model relation, in the course of natural language semantics. Thus, natural language

[20]Philosophers of language have thus long been indebted to Lewis Carroll (Carroll, 1960).

semantics does not do what is taken for granted by the logic in natural language thesis, and so, the thesis looks doubtful. There are some remaining issues, of course, such as whether we can press truth conditions into service to provide a consequence relation in some other, less direct way. I shall return to these below. But we can see from Lepore's argument that a model-theoretic consequence relation is not built into the basic apparatus of semantic theory.

To further support this conclusion, I shall explore a little more where model-theoretic techniques do fit into natural language semantics, and show how they do so without providing consequence relations.

II.1.2 Model Theory in Current Semantic Theory

Most all current work in semantics, including that work done in the 'model-theoretic' or Montagovian tradition, is in fact really doing absolute semantics. In this section, I shall illustrate this point by discussing a few of the important features of current model-theoretic semantics, and showing that they do not lead to consequence relations. Indeed, we will see that what is characteristic of such approaches to semantics these days is a reliance on *type theory*, not model theory in the sense needed to get logical consequence. Thus, we will see that the conclusion of the last section, that the basic idea of truth-conditional semantics does not lead to a model-theoretic consequence relation, is not specific to a neo-Davidsonian view of semantics. It follows just as much on a model-theoretic view. We will return in section II.3 to further questions of how model theory might find its way into absolute semantics.[21]

What is characteristic of most work in the model-theoretic tradition is the assignment of semantic values to all constituents of a sentence, usually by relying on an apparatus of types (cf. Chierchia and McConnell-Ginet, 2000; Heim and Kratzer, 1998). Thus, we find in model-theoretic semantics clauses like:[22]

(3) a. $[\![\text{Ann}]\!] = \text{Ann}$
    b. $[\![\text{smokes}]\!] = \lambda x \in D_e.\ x \text{ smokes}$

[21]This section is a brief discussion, focused on the issues at stake for the status of logical consequence in natural language. There are a great many more questions that are raised by the place of model-theoretic or other mathematical techniques in semantics. I address more of them in work in progress.

[22]In common notation, $[\![\alpha]\!]$ is the semantic value of α. I write $\lambda x \in D_e.\ \phi(x)$ for the function from the domain D_e of individuals to the domain of values of sentences (usually truth values).

We also find rules like function application (Heim and Kratzer, 1998; Klein and Sag, 1985):

(4) If α is a branching node and β, γ its daughters, then $[\![\alpha]\!] = [\![\beta]\!]([\![\gamma]\!])$ or vice-versa.

These are not things a neo-Davidsonian theory (one using traditional Tarskian apparatus) is going to have.

Even though clauses like this look different from those preferred by neo-Davidsonians, they provide absolute statements of facts about truth and reference. They just put those facts in terms of functions and arguments (as Frege would have as well!). We see that the value of 'Ann' is Ann, not relative to any model. The value of 'smokes' is a function, but one that selects the things that smoke, again, not relative to any model. Semantics needs to be absolute, but both model-theoretic and neo-Davidsonian semantic theories provide absolute truth conditions. These days, semantics in either tradition is absolute.[23]

Perhaps the most obvious of the distinctive features of the model-theoretic approaches is that it uses the typed λ-calculus to assign semantic values to all constituents. Compare the neo-Davidsonian (5a) (e.g. Larson and Segal, 1995) with its model-theoretic variant (5b):

(5) a. $Val(x, \text{smokes}) \iff x$ smokes
b. $[\![\text{smokes}]\!] = \lambda x \in D_e.\, x$ smokes

Though (5b) provides an object to be the semantic value where (5a) states that a relation holds, as far as truth conditions goes, these do pretty much the same thing, in the same way. That one posits a function is a difference in the apparatus the theories use (and so an ontological commitment in the long run), but not in the explanations they provide for the basic semantic property of this word. The speaker is not being attributed understanding of the theory of functions and relations, either for semantic values or Val relations. These are used to attribute to the speaker knowledge that 'smokes' applies to things that smoke.

Thus, the absolute nature of semantics, and the lack of appeal to a space of model that constitutes a consequence relation, is not specific to neo-Davidsonian semantic theories. It holds just as much for what is called

[23]Montague's original work (Montague, 1973) and subsequent presentations like Dowty *et al.* (1981) did officially rely on a notion of truth in a model. But even so, they usually drop reference to models when the linguistic analysis starts to get interesting.

the model-theoretic approach. That being said, there are some genuinely significant differences between model-theoretic and neo-Davidsonian theories, and the use of λ's is important for them. It is not my goal here to explore them deeply, but let me mention one, just to make clear that there are some, and they are important. The two sorts of theories disagree on the nature of semantic composition, and the use of λ's, really the use of a wide-ranging apparatus of functions, allows model-theoretic semantics to see semantic composition in terms of functions and arguments. Neo-Davidsonian theories see composition much differently, in many cases, in terms of conjunction (e.g. Pietroski, 2005). This does lead to some real empirical differences between model-theoretic and neo-Davidsonian analyses of various phenomena, and especially, has potentially far-reaching implications for issues of logical form in natural language. But this can happen while the two theories agree on the fundamental idea of specifying truth conditions as a way to capture speakers' linguistic competence.

I have dwelt at length on an internal dispute among semanticists about what sort of semantic theory works best, though in the end I have suggested that the dispute has been resolved, and everyone has opted for absolute semantics, in either model-theoretic or neo-Davidsonian guises. Current model-theoretic semantics does so using type theory, and so, where we have grown accustom to saying 'model-theoretic semantics' it might be better to say 'type-theoretic semantics'. Names aside, neither the apparatus of Tarski-style truth theories nor of type theory itself provides any sort of logical consequence relation, and in their uses in semantics they cannot. One way or another, semantics must be absolute, and so not relative to a model.

I suggest this makes vivid how the enterprise of truth-conditional semantics is distinct from that of studying logical consequence. Relative truth conditions, in Lepore's terminology, are just what anyone studying logical consequence should want. If we want to understand the logical properties of a sentence of a language, we look at how the values of the sentence can vary across models. This is just what the logics in formal languages thesis builds on. But semantics of natural language—the study of speakers' semantic competence—cannot look at that, and still capture what speakers understand. To capture what the speakers understand, semantics must be absolute, and so blind to what happens to a sentence across any non-trivial range of models. Thus, we cannot find the basic resources for studying logical consequence in natural language semantics, even in its truth-conditional form. We cannot take the step from the logics in formal languages thesis to the logic in natural language thesis. To give the argument of this section a name, let us call it the *argument from absolute semantics* against the logic in natural language thesis. The argument from absolute semantics, I submit, gives us good reason to reject the logic in natural language thesis.

II.2 Entailments and Consequences

The argument from absolute semantics shows that the semantics of natural language and the model theory of consequence relations are different things. But it might be objected that there are other ways to find consequence relations in natural language. Moreover, there is an obvious place to look to find such relations. Whether or not semantics is absolute, it must endow natural language with some implication properties. Indeed, implications are among the main data that empirical semantic theory builds on, and I doubt we would find any viable approach to semantics that failed to account for them. Even if the simple idea of capturing speakers' understanding of truth conditions does not hand us a consequence relation, it might be that the implications that are built into natural language do. If so, we would have an alternative route to defending the logic in natural language thesis. I shall argue in this section there is no such route. On the restrictive view of logical consequence we are now assuming, what we find in natural language are entailments, but not logical consequences.

As we saw above, sentences of natural language do present us with obvious implications, in that in certain cases competent speakers consistently judge that the truth of one sentence follows from the truth of another. Analytic entailments like the 'cousin' inference of section I.3 are good examples. And truth-conditional semantics is ready-made to capture some of these implications. It endows natural languages with *entailment* properties. Each sentence is associated with a set of truth conditions, and so truth-conditional containment properties are fixed. A sentence S entails a sentence T if the truth conditions of S are a subset of the truth conditions of T.[24]

We have already seen that in general, such entailment relations are not narrow logical consequence. Natural language provides entailments of other

[24] We might worry about whether extensional and intensional theories provide exactly the same entailment relations, and whether extensional theories really provide for strict entailment in the sense of Lewis and Langford (1932). But, I shall not dwell on that here, as both provide reasonable notions of entailment, and both turn out not to be logical consequence.

sorts. So, as I noted, advocates of the restrictive view of consequence will not find their preferred narrow notion in natural language entailments.[25] But if the only issue was that we find something more like the patterns of strict entailment in natural language, rather than narrow logical consequence, it would not impress anyone with even modestly permissive views. After all, the logic of such entailments, and many related notions, have been studied extensively.[26] At core, what would be needed is a somewhat different approach to the *necessity* constraint than classical logic uses, which would make possibilities more like metaphysically possible worlds than like classical models. Thus, it might be proposed, what we find in natural language may not be the most narrow of consequence relations, but it is close enough to what a restrictive view of consequence is after to offer an interesting version of the logic in natural language thesis.

I shall argue here this sort of response is not adequate. The reason is that natural language puts pressure on the *formality* constraint on logical consequence. Absolute semantic theory itself has no room for any such notion of formality. It specifies absolute meanings for each expression of a language, and sees no distinction between the ones whose meanings determine consequence relations and those which do not. And, I shall show, the facts about implication in natural language reflect this. Natural language is filled with rather idiosyncratic *lexical entailments*, driven by the meanings of a huge variety of lexical items. These depart too far from formality to satisfy restrictive views of logical consequence. In section III, I shall argue they depart too far for most permissive views as well.

We have already note that natural language provides us with non-logical implications, like analytic entailments. These, as the common wisdom goes, are determined by the meanings of non-logical terms like 'bachelor', not by the meanings of the logical constants. But natural language is permeated by entailments which strike us as evidently non-logical (by restrictive lights). Here is another case, much discussed by semanticists (see Anderson, 1971; Fillmore, 1968; Levin and Rappaport Hovav, 1995):

(6) a. We loaded the truck with hay.
ENTAILS
We loaded hay on the truck.

[25] As was observed also by Cresswell (1978).

[26] For some indications of the scope of this research, see among many sources Anderson and Belnap (1975), Anderson *et al.* (1992), Priest (2008), and Restall (2000).

b. We loaded hay on the truck.
DOES NOT ENTAIL
We loaded the truck with hay.

This is a report of semantic fact, revealed by judgments of speakers, both about truth values for the sentences, and about entailments themselves. It indicates something about the meaning of the word 'load' and how it combines with its arguments. More or less, the 'with' variant means we loaded the truck *completely* with hay, while the 'on' variant does not.[27]

To take one more much-discussed example, we see (Hale and Keyser, 1987; Higginbotham, 1989a):

(7) John cut the bread.
ENTAILS
The bread was cut with an instrument.

The meaning of 'cut', as opposed to e.g. 'tear', requires an instrument, as Hale and Keyser famously noted.[28]

Entailments like these are often called *lexical entailments*, as they are determined by the meanings of specific lexical items. As a technical matter, it is not easy to decide on the exact source of these entailments, i.e. whether they simply come from the atomic meanings of the verbs, the compositional semantics of verbs and arguments, or some form of predicate decomposition, presumably at work inside the lexicon. But this is an internal issue for natural language semantics (a very interesting one!). What is important for us is that any viable semantics must account for these sorts of entailment patterns. Though we have only looked at a couple of examples, I believe they make plausible the empirical claim that natural language is rife with such lexical entailments. Any semantic theory, of either the model-theoretic or neo-Davidsonian stripe, must capture them.

It is clear that these lexical entailments will not count as narrow logical consequences. The reason is not that we are looking at a wide space of classical models, rather than a smaller space like that of metaphysically possible worlds, or even some smaller space. Rather, the reason is that these entailments are fixed by aspects of the meanings words like 'load' and 'cut'. If we

[27]There is some debate about whether the connection of the 'with' variant to being completely loaded is entailment or implicature, but I believe the judgments on the (a/b) contrast above are stable enough to indicate an entailment.

[28]Hale and Keyser (1987) gloss the meaning of 'cut' as 'a linear separation of the material integrity of something by an agent using an instrument'.

start with any consequence relation which does not treat these as logical constants, we will not get the right entailments without violating the formality constraint.

From the restrictive point of view that we are adopting here, this shows that we will not find a logical consequence relation in the lexical entailments natural language provides. But, anticipating discussion of more permissive approaches to come, it is worth asking how far we would have to go to accommodate such lexical entailments in a consequence relation. Assuming we can take care of necessity, could we satisfy formality with an appropriately permissive notion of consequence which captures these cases? I shall offer two reasons we cannot. One is that we will find lexical entailments permeate language too far to generate what even modestly restrictive views would recognize as a logic. The other is that we will run afoul of another important condition on logical constants.[29]

Let us take up the first of these points first. Even if we were willing to consider taking e.g. 'load' and 'cut' to be logical constants, we would not have the 'logic of lexical entailments'. We obviously would not, as many more words than these trigger lexical entailments. Practically every one does! We cannot just take these two. We would have to make nearly every word a logical constant. This would render the formality constraint virtually trivial. The result would not be accepted as logic by restrictive views, but in trivializing the formality constraint, I doubt it would be accepted by most permissive views either.

There are a couple of reasons the situation might not be so dire, but I do not think they are enough to recover a viable consequence relation, according to restrictive or even moderately permissive views. It might not be that we have to take every individual word as a constant, as the patterns we are discussing occur across families of expressions. For instance, the pattern we see in (6) is one that reappears across a family of verbs, including 'spray', 'brush', 'hang', etc.[30] But as far as we know, we will still wind up with a very large and very quirky collection of families. Looking at the list of verb classes in Levin (1993), for instance, we find differences between classes of verbs of sound emission and light emission, and differences between classes of verbs

[29] Related points are made by Lycan (1989), though he emphasizes the difference in degree between logical consequences and lexical entailments more than I do.

[30] See Levin (1993) for many more examples. See Levin and Rappaport Hovav (2005) for a survey of some theories that seek to explain the sorts of entailments I am discussing here.

of bodily processes and verbs of gestures involving body parts. Taking each such class to indicate a logical constant would still give us a huge and quirky list of constants, which would still undercut the formality constraint. Even if it does not completely trivialize formality, it will still undercut formality well beyond what restrictive views can accept. I believe it would undercut formality enough to be unacceptable to modestly permissive views as well.

The only hope for narrowing down our domain of logical constants would be that underlying the many lexical entailments we find might be a small group of factors that generate them. In fact, it is a theoretically contentious issue just what generates the sorts of entailments we have been looking at, and what might group various word together into natural classes. But there is an optimistic idea that what appear to be idiosyncratic lexical entailments—entailments determined by the idiosyncratic properties of words' meanings—are really entailments determined by some hidden structure of classes of lexical items. But I suggest that as far as we know, we will still wind up with a group of 'logical constants' that are too quirky and too heterogeneous to satisfy formality in any substantial way. We will identify as formal a range of selected features driven by the quirky structure of natural language, not by the forms of valid arguments. For instance, it is a persistent thought that behind the entailments in (6) is something like a requirement for the 'with' variant that a container or surface be completely filled or covered.[31] (This constraint might have something to do with kinds of arguments of the verb, or just be coded up into the right meanings.) So our logic will have to be in part a logic of filling containers or covering surfaces. Restrictive views will reject this. And this is just the start. It will also have to be a logic of instruments for cutting to account for the entailments in (7). The result would be a formal listing of a quirky range of entailment properties, not the kind of logic the restrictive view is after. Again, it will not be the kind of logic modestly permissive views are after either, as I shall discuss more in section III.

This reminds us is that even though both logic and semantics are concerned with implications in general, they are concerned with different ones. Semantics is absolute, and interested in the specific meanings of all the terms in a given language. It is thus interested in the entailments that arise from those meanings. These are typically not logical consequence relations. The idiosyncrasy of natural languages makes this striking, as the 'load' case shows.

[31]See Dowty (1991) and Pinker (1989) for discussion and numerous references.

We find in natural language all kinds of unexpected and odd entailments, driven by the quirky meanings of various words. When we study these, we are not doing logic; we are certainly not looking at restrictive logical consequence.

There is a second reason for restrictive views to reject the idea of modifying logic to take in lexical entailments. By restrictive lights, the expressions we would have to count as 'logical constants' will not meet the standards for being logical constants. This will apply equally to the words like 'load', or to any possible underlying features that generate families of lexical entailments.

It is a common idea that to satisfy formality, the logical constants have to meet certain standards. The main standard is being 'topic-neutral', or not specifically about the particular things in the world, but only its general or 'formal' structure. This is indeed a further constraint than we have so far considered, but it is a natural way to fill out the idea of formality. In technical terms, it is often spelled out as a requiring that logical constants be invariant under permutations of the elements of the universe, or under the right isomorphisms.[32]

If we impose a constraint like this on logical constants, it is very unlikely that anything responsible for the kinds of lexical entailments we have been considering would meet it. Words like 'load' and 'cut' do not. Might some underlying features that might account for lexical entailments meet it? It is perhaps possible, but unlikely. One reason for skepticism is that it is well-known that practically no predicates and relations satisfy permutation invariance. For a domain M, the only predicates satisfying it are $\emptyset$ and M, and the only binary relations satisfying it are $\emptyset$, $M \times M$, identity, and non-identity (Peters and Westerståhl, 2006; Westerståhl, 1985). If we require normal permutation invariance as a constraint on logical constants, we will not find enough in the domain of predicates to capture the kinds of lexical entailments we are here considering. It might be that a more permissive view could appeal to a different range of isomorphisms to recover a notion of logical constant applying to some of the terms we have considered here, as I shall discuss more in section III. But nonetheless, this fact about permutation invariance shows how far from a restrictive view we would have to go to get a logical consequence relation from the lexical entailments of a natural language.[33]

[32]This idea is discussed by Mautner (1946) and Tarski (1986). See also van Benthem (1986) and Sher (1991).

[33]There has been some interesting recent work on conditions related to permutation

I have argued in this section that we will not find logical consequence relations in the lexical entailments of natural language. As we will need to refer back to this argument, let us cal it the *argument from lexical entailments*. This argument supplements the argument from absolute semantics. We will not find logical consequence relations in the basic apparatus of truth-conditional semantics, and we will not find it in the entailments present in natural language either. We thus have two reasons to reject the logic in natural language thesis.

II.3 Logical Constants in Natural Language

There is one more route I shall consider to saving the logic in natural language thesis. It might be objected that in spite of the points I made in the last two sections, there are some genuinely logical expressions in natural language, and their properties might somehow provide a consequence relation for natural language even if truth-conditional semantics and lexical entailments do not. In this section, I shall argue this is not so either. First, following up the discussion of model-theoretic semantics in section II.1.2, I shall show that the (model-theoretic) semantics of logical terms in natural language does not provide a genuine consequence relation. Second, I shall point out that natural language does not really come pre-equipped with a distinguished class of logical constants. If we wish to find logical constants in natural language, we have to identify them in ways that go beyond the semantics of natural language itself. This will point the way towards our discussion of how one can move from natural language to logic proper in section IV.

First, let us look at how the semantics of uncontroversially logical terms will be handled in an absolute semantics. As an example, I shall focus on the quantifiers, or more property the determiners: words like 'every', 'most', 'some', etc. Some of these are clearly logical constants by even the most restrictive lights, and aside from a few contentious cases, they satisfy the permutation invariance constraint we discussed in section II.2.[34] We will see that these logical expressions of natural language get an entirely standard model-theoretic semantic treatment, putting aside some structural differences

invariance (e.g. Bonnay, 2008; Feferman, 1999), but if anything, this work suggests even more restrictive criteria for logical constants. For a survey of related ideas about logical constants, see MacFarlane (2009).

[34] See Sher (1991) for a defense of the idea that generalized quantifiers are genuinely logical constants by restrictive measures.

between the most common logical formalisms and natural language; but they do so in a way that preserved the absolute nature of the semantics of natural language.

Methods of logic, particularly, of model theory, have proved extremely fruitful in the study of the determiners, as classic work of Barwise and Cooper (1981), Higginbotham and May (1981), and Keenan and Stavi (1986) has shown. Indeed, this led Higginbotham (1988) to declare that in natural language, model theory is the lexicography of the logical constants.

But model theory does this in a specific way. Consider, for example, a common way of representing the meaning of a determiner like 'most':

(8) $[\![\text{most}]\!](A, B) \iff |A \cap B| > |A \setminus B|$

The meanings of determiners, on this view, are relations between sets expressing cardinality properties. This definition is drawn from the theory of generalized quantifiers. This is a is a rich mathematical field, and it has led to an impressive number of non-trivial empirical predictions and generalizations in semantics.[35]

The application of generalized quantifier theory to the semantics of determiners is one of the most well-explored applications of model theory to natural language, and a key example of the success of model-theoretic semantics. But it remains an example of *absolute* semantics, as the arguments of section II.1 show it must. There is no tacit quantification over a domain of models in the semantics of determiners. The sets A and B in the above definition (8) must be drawn from whatever fixed domain of individuals is involved in the rest of the semantics. There is no tacit 'for any model $\mathfrak{M}$'. In this way, we can do what is often labeled 'model theory', while still doing absolute semantics, and still not generate a consequence relation.

Actually, this point is already enshrined in the theory of generalized quantifiers, as the distinction between *local* and *global* generalized quantifiers.

(9) a. Local: $[\![\text{most}]\!]_M = \{\langle A, B\rangle \subseteq M^2 \colon |A \cap B| > |A \setminus B|\}$
b. Global: function from M to $[\![\text{most}]\!]_M$

The direct application of generalized quantifiers to semantic theory uses *local* generalized quantifiers, as an absolute semantics should.

[35]The mathematics of generalized quantifiers was first studied by Lindström (1966) and Mostowski (1957). See Barwise and Feferman (1985) and Westerståhl (1989) for surveys. Among the most striking empirical predictions is the *conservativity* constraint articulated by Barwise and Cooper (1981) and Higginbotham and May (1981). This may well be a semantic universal.

On the other hand, if we are to embed generalized quantifier theory in our theory of logical consequence, it is global quantifiers that we need. As has been much discussed, to capture logical consequence relations with quantifiers, the domain must be allowed to vary. That is just what global generalized quantifiers do. (We will return to this issue in section IV.) For the study of semantics of natural language—absolute semantics—local generalized quantifiers are the basic notion; while for the study of logical consequence, global ones are.

Attending to the local versus global distinction, we can reconcile two facts that might have seemed in tension. First, familiar determiners in natural language have more or less the semantics that logical theory says they should. Though there is some interesting linguistics subtlety (a little of which I shall mention below), 'every' is pretty much a universal quantifier as we come to learn about in logic. And yet, the semantics of this expression is absolute, and does not make reference to the range of models essential to the logic of quantification. But the reason is simply that semantics of natural language only uses local properties of quantifiers in spelling out the semantics of determiners. These are readily available for absolute semantics.

I thus conclude that the presence of recognizably logical expressions in natural language doe not help support the logic in natural language thesis. We can find terms which we recognize as logical, and give them essentially the semantics logic should make us expect, while keeping semantics entirely absolute, and not involving any true consequence relations.

Now, this does not mean we can never look at the global notion of quantifier in thinking about natural language. The basic idea for giving absolute truth conditions is the local one, and in fact, sometimes we can get interesting further results out of local definitions.[36] But on occasion, we learn something by abstracting away from absolute truth conditions, by looking at global generalized quantifiers. A example is the idea that natural language determiners express restricted quantification. This is captured in two ways:

(10) a. CONSERV (local): For every $A, B \subseteq M$, $Q_M(A, B) \iff Q_M(A, B \cap A)$

b. UNIV (global): For each M and $A, B \subseteq M$, $Q_M(A, B) \iff$

[36] For instance, counting and classifying available denotations over a fixed domain can be interesting. One example is the 'finite effability theorem' of Keenan and Stavi (1986), which shows that over a fixed finite universe, every conservative type $\langle 1, 1\rangle$ generalized quantifier is the denotation of some possibly complex English determiner expression.

$$Q_A(A, A \cap B)$$

UNIV is a generally stronger principle (Westerståhl, 1985), and captures an interesting way in which natural language determiners really quantify only over a restricted domain. If we refused to look at global properties of generalized quantifiers, we would not see it.

In looking at this sort of global property, we are not simply spelling out the semantics of a language. Rather, we are abstracting away from the semantics proper—the specification of contributions to truth conditions—to look at a more abstract property of an expression. It turns out, in this case, abstracting away from the universe of discourse is the right thing to do. Particularly when asking about logical or more generally mathematical properties of expressions, this sort of abstraction can be of great interest. And, we can prove that typical natural language determiners satisfy UNIV, invoking a little bit of mathematics, even if it goes beyond the semantics of any language per se.

This sort of possibility shows how we might take the step from semantics proper to logic, as I shall discuss more in section IV. But for the moment, I shall pause to note one feature of how the application of the model theory of generalized quantifiers to natural language works. Rather than indicating anything like a genuine consequence relation in natural language, it illustrates how the application of model-theoretic techniques to natural language semantics really turns out to be an instance of the general application of mathematical techniques. In applying local generalized quantifiers, what we are really applying is a little bit of set theory (deeply embedded in standard model theory too), to represented cardinality-comparing operations expressed in language, and study their properties. We can do so over an arbitrary or a fixed domain, depending on the question under investigation. But especially when we are stating absolute truth conditions, it is the mathematics of cardinality comparison over a fixed domain that we really invoke. Though this can reasonably be counted as logic, as it is the subject-matter of local generalized quantifiers, it is not *logic* in the core sense of logical consequence relations. When we use logic in the study of natural language semantics, we are typically using logic in the broad sense in which logic can be found in many domains, not the narrow one of the study of logical consequence relations.[37]

[37]The application of mathematical techniques to natural language semantics is not specific to techniques found in model theory, nor is it specific to the logical expressions in

The case of determiners shows how we can grant that there are logical expressions in natural language, which get more or less standard logical analyses, and still reject the logic in natural language thesis. But, one might still wonder if natural language supports the logic in natural language thesis in another way, by presenting a distinguished class of logical terms. Even if their semantics is absolute, this might be a partial vindication of the thesis.

We have already seen some reason to be skeptical of whether language does the job of identifying the logical constants for us. We have seen that the mere fact that some expressions are well-analyzed by techniques like those of generalized quantifier theory only shows that they are amenable to a kind of mathematical analysis. In fact, the class of expressions that are at least partially analyzable in mathematical terms is very wide, and contains clearly non-logical expressions, as well as plausibly logical ones. To cite one example, a little bit of the mathematical structure of orderings (a tiny bit of topology) has proved useful in analyzing certain adjectives (see Kennedy and McNally, 2005). As we discussed in section II.2, there are good reasons to doubt expressions like these are logical, and they are certainly not by restrictive lights. So, natural language will not hand us a category of logical constants identified by having a certain sort of mathematically specifiable semantics.

Is there anything else about a language—anything about its grammar, semantics, etc.—that would distinguish the logical constants from other expressions? No. Generally, expressions we are inclined to accept as logical constants by restrictive lights, like quantifiers, conjunction, negation, etc., group with a much larger family of elements of a natural language. They are what are known as 'functional' categories, which include also complementizers ('that', 'which'), tense, mood, and other inflectional elements, morphological elements like degree terms ('-er'), etc. By strict restrictive standards, it is doubtful that all of these will be counted as logical constants.[38] Some of these elements, like tense, have been points of dispute among restrictive views over what counts as logical. But most restrictive views will refuse to count the complementizer 'that' or the comparative morpheme '-er' as logical constants.

This claim relies on the highly restrictive tendencies of traditional re-

natural language. I discuss more ways that mathematics applies to absolute semantics in work in progress.

[38]As we will see in section IV.3, even the elements we are inclined to accept as logical do not behave exactly as the expressions in formal languages do.

strictive views, which simply find no place in their logics for such terms. Whether or not they fail to meet some criterion, like permutation invariance, is a rather more delicate matter, and I will not try to decide it here. One of the typical features of functional categories is that their semantics is not the familiar semantics of predicates and terms; rather, their semantics tends to be 'structural', involving times or worlds for inflectional elements, degrees for comparatives, etc.[39] We cannot easily apply permutation invariance to these, and instead, we will have to ask if they satisfy the right property of invariance under isomorphism. Whether they do will depend on just what isomorphisms are at issue.[40] This will become more of an issue when we turn to permissive views in section III. But as with other abstract properties of expressions, it is not something that is specifically marked by the grammar. Natural languages do not tag certain expressions as invariant under isomorphisms—absolute semantics cannot do this! Rather, as we saw with UNIV, they provide absolute semantics for expressions, and we can work out mathematically what would happen if we take a global perspective and look at invariance properties. Hence, if we are to conclude that some or all functional categories meet some invariance conditions, we must go beyond what the grammars of our languages tell us.

Linguistically, there are lots of distinguishing marks of the class of functional categories. Functional categories are closed classes. Unlike nouns, verbs, and adjectives, it is impossible to add to these classes (except at the glacial pace of language change). You cannot simply add a new determiner or a new tense to your language the way you can easily add a new verb or noun. Functional categories also play special grammatical roles. In many languages some determiner, even a semantically minimal one, is needed to make a noun phrase an argument of a predicate. Generally, functional categories act as 'grammatical glue' that binds sentences together. This sort of grammatical role is clearly on display with complementizers. In a more theoretical vein, unlike nouns, verbs, and adjectives, functional categories do not assign argument structure.[41]

But, these sorts of features group all the functional categories together. They do not distinguish the logical sub-class from the wider one, as they do

[39] Many linguists will gloss this as their not having a 'theta-involving semantics', meaning it does not involve regular predicates and their linguistically provided arguments.

[40] Hence, we do find, for instance, logics of comparison, as in Casari (1987).

[41] Many recent syntax texts will review the notion of functional category. See also Abney (1987), Fukui (1995), Grimshaw (2005), and Speas (1990).

not really pertain to logicality at all. Indeed, I do not know of any linguistic property that does distinguish logical elements from other functional ones (assuming the restrictive view that they are not all logical). So, I conclude, natural language does not sort expressions into logical and non-logical. It contains expressions that will count as logical—just how many depends on your standards for logical terms—but it does not itself distinguish them. That is something we can do when we look at a language from a more abstract perspective, as we will discuss more in section IV.[42]

We thus see one more way that the logic in natural language thesis fails. Let us call this one the *argument from logical constants*. It joins the argument from absolute semantics and the argument from lexical entailments in showing what is wrong with the logic in natural language thesis. In particular, like the argument from lexical entailments, it shows that a route to finding significant logical structure in natural language is not open.

This concludes my main discussion of the logic in natural language thesis. But there is one loose end to tie up, and one further issue to address. Both return to points we have seen in passing in this section. First, the conclusions I have reached have been under the assumption of a restrictive view of logical consequence, and we need to see how they fare with a permissive view. We will do this in the next section III. Second, we have already seen ways that we can go beyond semantics proper to explore logical properties of natural language expressions. This reminds us that we can take the step from natural language to logic, even if the logic in natural language thesis is false. We will explore how this step may be taken in section IV.

III The Logic in Natural Language Thesis from the Permissive Point of View

Above I offered three arguments against the logic in natural language thesis, but generally assumed a restrictive view of logical consequence. In this section, I shall examine how those arguments fare if we adopt a permissive view, as I did at a few points above as well. I shall argue that we still have some good reasons to reject the logic in natural language thesis. Even so,

[42]This conclusion is not so far off from the claim of Evans (1976) that what he calls 'semantic structure' is distinct from 'logical form', and does not determine logical consequence.

permissive views will have more opportunities to find logical properties in natural language. I shall suggest that this will not really support the logic in natural language thesis, without also weakening it to a point where it threatens to become uninteresting. But at the same time, as we will discuss more in section IV, it will show how permissive views might see logic and natural language as more closely linked than restrictive views do.

The three arguments I offered against the logic in natural language thesis in section II were the argument from absolute semantics, the argument from lexical entailments, and the argument from logical constants. The argument from absolute semantics was the main argument, which showed that a natural language does not have a consequence relation in virtue of having a semantics. The other two arguments served a supporting role, by showing that we will not find restrictive logical consequence in natural language by other means.

The argument from absolute semantics in fact made no use of the premise of restrictive logical consequence. It made no assumptions about logical consequence beyond that it will require a non-trivial space of models (with more than one element). That is not something even the permissive views in question will deny. So, I conclude, the argument from absolute semantics is sustained on a permissive view. As this was my main reason for rejecting the logic in natural language thesis, I believe that thesis looks doubtful even on a permissive view of logical consequence.

We have already seen in sections II.2 and II.3 that matters are somewhat more complicated with the other two arguments. Let us consider the argument from lexical entailments first. The idea was that one might find a logical consequence relation in the entailments presented in natural language, and use that as a route to the logic in natural language thesis. We saw in section II.2 that those entailments are lexical entailments, and fail to be logical consequences according to restrictive requirements. But we also saw that there might be ways to count some of them as logical by more permissive lights. I argued in section II.2 that we could not do so without undercutting the formality condition. I offered two reasons for this. One reason was that we would have to accept logical constants which would not meet the permutation invariance condition. The other was that the needed logical constants would permeate language so extensively that we could only preserve the logic in natural language thesis at the cost of weakening formality too much. In the worst case, preserving the logic in natural language thesis threatened to trivialize formality.

Of course, extremely permissive views can reject both these claims. The question is what the results of doing so would be. Let us consider permutation invariance first. Though this could be rejected out of hand, it might make the formality constraint unacceptably weak if there is are no restrictions on what counts as a logical constant. So, the more likely avenue for a permissive approach, as we saw above, is to find some alternative constraints. The most obvious way, as we also saw, is to look for some appropriate structure for which a candidate expression will be invariant under isomorphism. Returning to the 'load' example (6), for instance, we might wonder if we can find an appropriate structure for filling containers or covering surfaces, which would render it invariant. A hint of how that might be done is its similarity to ideas from mereology. So, though I am not sure just how it would go, there is no reason a permissive view might not be able to find a logic of filling and covering, and capture the 'load' entailments through it. More generally, case by case, a permissive view might be able to see each lexical entailment as derived from some sort of logical constant.[43]

For each individual case, this will be plausible enough by permissive lights. After all, permissive views are interested in finding many different candidate logics, and will not find it odd to look for a logic of e.g. filling or covering. There is some risk of losing any sense of what makes a logical constant a logical constant, but presumably good permissive views will find some way to keep this notion reasonably constrained. The more pressing worry is about how pervasive the lexical entailments of natural language are, and what accounting for all of them as logical would do to formality.

At the extreme, as we saw in section II.2, to sustain the logic in natural language thesis we might need to count nearly every word in a language as a logical constant, as nearly every expression will generate some lexical entailments. Even if we are willing to grant that many of them can be viewed as permissive logical constants, making nearly all of them constants will radically weaken formality. If nearly every expression is a logical constant, then there is little left of the idea that inferences are valid based on distinguished formal structure. That might rescue the logic in natural language thesis, but at the cost of trivializing it.

The more hopeful option for the permissive defense of the logic in natural language thesis is that there might be a smaller set of features that generates the lexical entailments. As I discussed in section II.2, it is not clear whether

[43]As I mentioned in footnote 33, there has been some exploration of alternatives to isomorphism invariance, but they have tended to offer more, not less, restrictive conditions.

this is so, and if it is, it is not clear just how large or how heterogeneous the features involved will be. Again, a permissive view will likely be willing to grant that each such feature individually can be a logical constant, and have a logic for the entailments associated with it. The question, again, is what we get if we put them all together. Here, if indeed the set is small enough and homogeneous enough, there could possibly be a defensible permissive view. But I am doubtful that will work. If the set is too large, it will trivialize the notion of formality again. If it is too heterogeneous, it will undermine the logic in natural language thesis in other ways. If our group of logical constants is too heterogeneous, it is not clear if we will find any single consequence relation which makes sense of all of them in one logic, as the logic in natural language thesis demands. Even if we can, it is not clear if the result will wind up being any better than a brute force way to capture the logic of natural language by enumerating the lexical entailments. Technically, that would not falsify the logic in natural language thesis, but it would undercut its interest. The thesis was interesting if we had some, even permissive, idea what counts as logic, and thought we could find this logic in natural language. A brute force coding of lexical entailments would not provide that. Of course, finding any tractable way to enumerate all the lexical entailments in natural language would be of huge interest! The problem is that it may not produce a natural or interesting consequence relation, which is what the logic in natural language thesis looks for.

I thus grant that the argument from lexical entailments might not work on a permissive view (or at least, an extreme permissive view). Even so, I still register a great deal of skepticism over whether the argument can be bypassed in a way that leaves the logic in natural language thesis substantial and interesting. As the argument from absolute semantics still holds, I think we should be dubious of the logic in natural language thesis by permissive lights, and I count my skepticism about ways permissive views might bypass the argument from lexical entailments as more reason to doubt the logic in natural language thesis. Yet all the same, we should grant that by permissive lights, the relation of lexical entailment to logic is not so clear-cut.

Finally, we need to reconsider the argument from logical constants from section II.3. This argument shows that natural language does not present us with a group of distinguished logical constants, thus blocking one other potential route to the logic in natural language thesis. Part of this argument goes through for permissive views. As the semantics of all terms, logical ones included, must be absolute, we will not find any terms of a natural language distinguished within the language by having any particular global or non-absolute logical properties. The semantics and the rest of the grammar of a natural language does not do that. This part of the argument does not rely on restrictive assumptions, and works equally well for permissive views.

But there is another part of the argument from logical constants that is likely to fail by permissive lights. I argued that the only grammatically distinguished class of terms that contains the uncontroversially logical ones is the class of functional expressions. By restrictive standards, this proved too large to count as the class of logical constants. But it remains open to permissive views to simply accept that all functional expressions are logical constants. As we discussed in section II.3, it will be easier to sustain this claim than the corresponding one for lexical entailments, as functional categories do have features that may make them amenable to logical analysis. Permissive views hunting for new interesting logics might well find the functional categories of natural language fertile ground.

Some of the worries I just raised surrounding the argument from lexical entailments will apply here as well. It is not clear if the result of counting all functional expressions as logical constants will be a single coherent consequence relation, or a brute force coding of multiple logics corresponding to multiple classes of functional elements. Hence, as with lexical entailments, I remain skeptical about whether permissive views of logical constants will offer a viable route to the logic in natural language thesis, even if they are able to undermine some parts of the argument from logical constants of section II.3. All the same, I grant, this is one of the more promising possibilities for permissive views of logical consequence to consider.

I conclude that for permissive views, the logic in natural language thesis still looks doubtful. The argument from absolute semantics still holds, and there is good reason to be skeptical of the possibility of avoiding the arguments from lexical entailments and logical constants. Both these latter two arguments still have some force, even if they are weakened on permissive views. Yet the relative weakness of the arguments from lexical entailments and logical constants does indicate possibilities for permissive views to explore, which might lead to something related to the logic in natural language thesis, I suspect in a much-weakened form.

It is clear that permissive views of logic can readily find properties of natural language that lead to interesting logics (as they can in many other domains as well). This does not by itself sustain the logic in natural language thesis, but it reminds us that both permissive and restrictive views can find logically relevant material in natural language. Not surprisingly, permissive views do so more easily, but both can. How this can happen, without the logic in natural language thesis, will be the topic of the next section.

IV From Natural Language to Logic

Let us suppose the logic in natural language thesis is false. This claim may be more secure by restrictive than permissive lights, but I have given some reasons to think so for both views. Regardless, we have also seen throughout the discussion above that we can often find things of interest to logic in natural language. We have seen that we can find expressions in natural language which turn out to have interesting logical properties, or even turn out to be logical constants, and we can find entailments which might prove of interest to permissive views of logic. This should not be great surprise, as it has been clear all along that natural language does present us with a range of implication phenomena, and range of expressions which might be of interest to logic. We can get from language to logic somehow, and this is a fact that logicians of both permissive and restrictive varieties have often exploited.

The question I shall explore in this section is how the step from natural language to logic can be taken. I shall argue that the space between natural language and logic can be bridged in a fairly familiar way, by a process of *identification* of logical constants and *abstraction*. But I shall highlight that we also need a third component of *idealization*, and argue that component is more substantial than many logicians might have assumed. Together, these processes allow us to move from the absolute semantics of natural language proper to a logical consequence relation. Isolating them will also help us to see how much is required for permissive or restrictive views to make the jump from natural language to logic.

A metaphor might help illustrate what the processes do. The richness of natural language shows us logical consequence 'in the wild', in our discursive and inferential abilities and practices. Sometimes we want to take it back to the lab, magnify some part of it, dissect it, modify it, and purify it. What we get is a cousin of the wild type, just as implications can be distant cousins of narrow logical consequence. How distant a cousin, of course, depends on what your view of consequence was to begin with. Frege did something like this with quantifiers. The problem of polyadic quantification theory

Frege solved came up in part because it came up in natural language (and in mathematics as well). But neither Frege's solution to the problem, nor many model-theoretic variants, amount to doing natural language semantics. Rather, they require identification of the relevant logical features of polyadic quantification, abstraction from the meanings of non-logical expressions, and idealization from features of natural language grammar not relevant to the core logical issues. That is what we generally do when we move from natural language to logic.

I shall address each of these processes in turn. Actually, I shall start out of order by discussing abstraction, as it is the main process by which we move from the absolute semantics of natural language to model-theoretic consequence relations. I shall then discuss the role played by identification, especially when it comes to logical constants, even though identification typically has to happen before abstraction. Finally, I shall discuss the important role for idealization in getting us the kinds of logics that have proved so useful.

IV.1 Abstraction from Absolute Semantics

I have argued that the semantics of natural language is absolute, and so does not provide a consequence relation. But there is a way to move from absolute semantics to the kinds of models we need to build model-theoretic consequence relations. Once we have them, we will be well-positioned to start working out logics, by looking at (the right sort of) preservation of truth or designated value across the range of models we produced. Presumably the resulting consequence relations will be genuinely logical, and also reflect some aspects of the absolute semantics we started with. That would be a major step towards building a laboratory-refined logic out of a wild-type natural language.

Actually, the way to do this is well-known from discussions of logical consequence and model theory. All we need to do is abstract away from the absolute features of our semantics. Absolute semantics will typically provide extensions for predicates, and referents for terms, etc. It will provide more than that, as it will provide semantic values for many more sorts of natural language expressions. It might also provide structures for tense, mood, etc., which might enrich the structures we are working with. But for now, let us simply focus on familiar extensions for predicates and terms.

How can we get from absolute extensions to a space of models? As has been much-discussed, we need to somehow abstract away from the specific meanings of the expressions of a language. The obvious way to do this, if we are thinking of extensions as sets of individuals, is simply to allow these sets to vary. How much they are allowed to vary will determine properties of the logic that results, but the basic idea is that they should vary freely. We might also approach abstraction by allowing substitutions of different expressions with different meanings (as discussed e.g. by Quine 1959, 1986, following Bolzano), but with enough set theory around, simply allowing the sets to vary is the natural way to abstract from the specific meanings of terms, and it avoids some problems that are well-known for substitutional approaches. Doing this will get us a non-trivial space of many models.

It is well-known, however, that just varying the extensions of predicates and terms is not enough. Only doing this will make our consequence relation dependent on what actually exists in the world (or in the domain of our absolute semantics). Many philosophers have concluded this is not enough to satisfy necessity, and not enough to satisfy the topic-neutrality idea behind formality. But it is also well-known what is needed to do better: we need to let the domain from which we draw individuals (the 'universe of discourse') vary as well, and draw our extensions for predicates and referents of terms from those varying domains.[44] Varying the domain, and the extensions of predicates and terms, produces a wider space of models. Indeed, as described, it produces the space of models that generates classical logic. But the main idea here is quite general, and can be applied to generate different sorts of consequence relations, depending on what structure is involved in the models, and how variation is understood. The case of classical logic gives us good reason to think that this process of varying domains and extensions is a successful one for getting us from absolute semantics to the kinds of spaces of models that serve logical consequence.

This is the process I shall call *abstraction*, as it involves abstraction from the meanings of terms. Abstraction gets us from absolute semantics to a space of models, and so, get us the essentials of a logical consequence relation. I understand abstraction as a *kind* of process, that can be implemented in different ways, resulting in different logics.

[44]This is the received view in the tradition stemming from Tarski. There are some historical issues about just where and how domain variation was first recognized, but I shall not pursue them here. See the references mentioned in footnote 6.

IV.2 The Identification of Logical Constants

The use of abstraction to move from absolute semantics to logical consequence is standard, at least in the post-Tarskian tradition. But, as is very well-known, it also brings us back to the issue of logical constants. We cannot abstract away from the meaning of every expression of a language and get back an interesting consequence relation, or one that will count as a consequence relation by restrictive lights. We have to keep fixed the meanings of the logical terms, and that means identifying the logical constants. Indeed, we have to identify the logical constants before we can abstract away from the meanings of the non-logical terms. The formality constraint requires that we identify the right formal structure for valid inferences, and the logical constants provide that structure.

We discussed the issues of how to get logical constants from natural language in section II.3 and in section III. That discussion reminds us that when we abstract away from absolute semantics, we will need to make sure our logical constants get suitable *global* interpretations in our space of models, or are properly handled when we define satisfaction in a model. But the more important result of those sections, for our current concern, is that natural language does not do the job of identifying the logical constants for us. At least, if we are not so permissive as to count virtually every term as logical, or perhaps every functional category, then natural language will not distinguish the logical constants for us.

Thus, if we are to carry out the process of abstraction to get a consequence relation, we will also have to carry out a process of *identification* to identify the logical constants. I argued above that natural language does not do the identification for us, and we will have to do it our selves. The discussion of section II.3 gave a few indications of how we might proceed. By restrictive lights, at least, we might appeal to some property like permutation invariance, or some global property of isomorphism invariance. (Such global properties might have to be articulated together with the abstraction process that would give us a space of models in which isomorphism could be defined.) If we are very permissive, we might find the process easier, and perhaps more a matter of identifying terms of interest than terms that meet some restrictive standard. But nonetheless, aside from the extreme case, we will have to do the identification, be it easy or hard to do. Doing so goes beyond what natural language itself tells us. We will have to go beyond observations of natural language 'in the wild', to our laboratory setting, and do the work of identifying the logical constants ourselves.

Abstraction and identification work together, and in doing so, they characterize a common post-Tarskian understanding of model-theoretic consequence relations. Isolating them helps to make clear how we can get to such consequence relations starting with absolute semantics of natural language. They also allow us to measure different ways that more or less permissive views might carry out the processes differently. We have seen that some permissive views might find the process of identification easier. Some related views might find the process of abstraction harder, as, for instance, constraints on metaphysical possibility or the structure of strict entailment might be required. Especially when it comes to identification, we will have to rely on some prior grasp of what logic is about to carry out the process. Perhaps, as we have been assuming, notions like formality or permutation invariance will be our guide. Regardless, both processes mark ways in which we depart from natural language when we build consequence relations.

IV.3 The Idealization Problem

The two processes of identification of the logical constants and abstraction from lexical meaning give us the tools to get from natural language to logic, and their use in tandem is quite familiar. But, I shall argue here, they are not sufficient to get us something we would want to call a logic. One more step, *idealization*, is needed. The reason is that even after we have performed abstraction, we are still going to be stuck with some idiosyncratic and quirky features of natural language grammar, that we will not want to contaminate our logic. Even after we have abstracted away from absolute lexical meaning, we still have not finished our laboratory work, and further step of purification will be in order.

To illustrate the kinds of features of natural language we will want to idealize away from, let us return to the behavior of quantifiers in natural language. Even when we have recognizably logical expression, like ‘every’, their behavior in natural language will be determined by a wide range of features of the grammar. This will not undercut the expressions having the kinds of semantics logic leads us to expect (in local form), but it can, and does, produce some unexpected behavior, that make them not work quite the way we expect a logical constant figuring in a consequence relation to work.

We see this, for instance, in the scoping behavior of the determiner ‘every’.

'Every' is a universal quantifier, no doubt, and behaves like one. In fact, it is a *distributive* universal quantifier, in a way that 'all' is not in English. As we have seen, it gets essentially the semantics of universal quantification, in local form. But natural language quantifier scoping behavior is well-known to be complex and delicate (or if you like, very quirky), compared to the essentially uniform scope potentials most every formal language builds in. To start with, there are well-known subject/object asymmetries in scoping which affect 'every'. They are brought out most strikingly if we also throw in a negation for good measure. Observe the following, as judged by native speakers, is *not* ambiguous:

(11) John did not ready every book.

This sentence does not have a reading where 'every' scopes over the negation, i.e. it cannot mean $\forall x(B(x) \rightarrow \neg R(j,x))$. On the other hand, if we substitute in an existential (i.e. an 'indefinite'), then the reading is available. The following has just the scope ambiguity standard logic would lead you to expect:

(12) John did not read a book.

In fact, this just scratches the surface of the maddening behavior of quantifier-like constructions in natural language. Even so, it is enough to show how unlike the nice uniform behavior of quantifiers in formal languages the behavior of quantifiers in natural language can be.[45]

Existential quantifiers in natural language show other sorts of behavior that we might not expect, if logic were to be our guide. Not least of which, existential and universal quantifiers do not share the same distribution in some puzzling environments, such as:

(13) a. i. Max is a friend of mine.
ii. * Max is every friend of mine.
b. i. There is a book on the table.
ii. * There is every book on the table.

Indeed, if we thought 'there is' expressed existential quantification, or 'heralded existence' in the much-noted phrase of Quine (1958), we might already

[45]The literature on these issues is large, but for some starting points, see Aoun and Li (1993) on subject/object asymmetries, Beghelli and Stowell (1997) on 'every', and Fodor and Sag (1982) on indefinites. The latter has spawned an especially huge literature.

be puzzled as to why it can pair with a quantifier at all. But even putting that aside, nothing in our familiar logical formalisms tells us to expect that some quantifiers occur in some positions, and some in others, or that existential (and related) quantifiers show special behavior.[46]

These are two illustrations, specific to quantifiers, of how the behavior of logical constants in natural language can be quirky. More specific to our concerns is that presumably we do not want our notion of logical consequence to include such quirks. I take it we do not want a notion of logical consequence captured in a language that builds in significant distributional differences between quantifiers, or builds in different scope behavior for existential and universal quantifiers.

Of course, when we build formal languages of the usual sorts, we smooth out this sort of quirky behavior. Standard syntax for formal languages provides the same distribution for elements of each category, so all quantifiers enjoy the same distribution. We thus will not have the kind of behavior we see in (13). When it comes to scope, standard formal languages mark scope explicitly, by such devices as parentheses. Thus, the kinds of questions about the possible scopes of quantifiers in surface forms we encounter in (11) and (12) will not arise once we fix the syntax of scope in a formal language. My point is that when we set up formal languages this way, we are not simply reflecting the grammar of the logical constants in the natural languages with which we start. We make substantial departures from the structure of natural language when we set up familiar sorts of formal languages.

We do so for good reason. We want our formal languages to display uniform grammatical properties in important logical categories. We want our quantifiers to show the important logical properties of quantifiers, for instance, not the quirky properties they pick up from the grammar of natural language. I am not going so far as to claim that any formal structure infused with such natural-language peculiarities would fail to be a logic (by permissive lights!). At least, those with permissive views might be interested in such formal languages, and some logical apparatus has been used to address e.g. quantifier scope in natural language.[47] But when we think about the core idea of logic as the study of valid arguments, it is just hard to see why these sorts of quirks should be relevant. Good reasoning about 'every' versus

[46]Again, the literature on these issues is large, but the phenomenon illustrated in (13) was first discussed by Milsark (1974).

[47]For instance, see the textbook discussion of Carpenter (1997), or the very interesting work of Barker (2002).

'some' just does not seem to be different in the ways grammar sometimes makes these words behave differently. So, in building logical consequence relations that reflect this core idea, we should *idealize* away from these quirks of grammar, even if we see them in natural languages. It is clear that restrictive views will have to idealize this way. In fact, even very permissive approaches to logic usually idealize in just the same way, when they are exploring core logical consequence. We see this in the standard formalisms of non-classical logics (e.g. Priest, 2008). Both permissive and restrictive views will typically idealize substantially in moving from natural language to logic.[48]

Idealization, as it figures here, is a familiar kind of idealization in scientific theorizing, that builds idealized models. One way to build idealized models is to remove irrelevant features of some phenomenon, and replace them with uniform or simplified features. A model of a planetary system is such an idealized model: it ignores thermodynamic properties, ignores the presence of comets and asteroids, and treats planets as ideal spheres (cf. Frigg and Hartmann, 2009). When we build a logic from a natural language, I suggest, we do just this. We ignore irrelevant features of grammar, and replace them with uniform and simplified logical categories. We do so for particular purposes. If we want to think of logic as the general, and highly idealized, study of valid arguments, such an idealized tool will serve our purposes well. But different purposes will require different tools. If we want to use logic as a tool for analyzing the structure of natural language in detail, as I mentioned, different idealizations will be in order. For the kinds of concerns with core logical consequence we have been focusing on here, we will want to idealize away from the quirks of natural language grammar; and regardless, some idealization will be in order for any purpose.

Thus, we need to add a process of idealization to those of abstraction and identification. We need all three to get from natural language to logic. We only get to logic—something that serves our purposes in analyzing valid reasoning, and is recognizably like what logicians work with—when we include idealization.

[48]Quine (1960) strikes a similar note in his discussion of the role of 'simplification' in regimentation.

V Conclusion

I have argued for two main claims in this paper. First, the logic in natural language thesis is false: we do not find logical consequence relations in our natural languages. Though the logics in formal languages thesis might have made the logic in natural language thesis seem appealing, it still should be rejected. Part of my argument for this claim relied on a restrictive view of logical consequence; but only part did, and we saw several reasons to doubt the logic in natural language thesis on permissive views as well.

Second, I have tried to show how natural language can serve as a useful object of study for those of us interested in logic, even if the logic in natural language thesis is false. The history of logic tells us that this must be possible, as many advances in logic have come from at least glancing at our own languages. As the quote from Beall and van Fraassen above reminds us, arguments are presented in language, so we have little chance of producing an analysis of valid arguments which paid no attention to language all. Indeed, as we considered the semantics and grammar of natural language above, we found many things of great interest to logic. To explain this, in light of the failure of the logic in natural language thesis, I sketched a three-fold process that allows us to get from a natural language to a logical consequence relation. The process involves identifying logical constants, abstracting away from other features of meaning, and idealizing away from quirks in the structure of human languages. The relation between logic and natural language is thus less close than advocates of the logic in natural language thesis would have it, while the three-fold process allows that there is some connection.

As I said at the outset, this is an autonomy thesis. Natural language semantics per se does not include logic. Logic, in its core, does not include the quirks or specific contents of natural language semantics. Moreover, to get from natural language to logic, you will have to rely on some fairly robust, if perhaps general, ideas of what makes something count as logical. You cannot find logical consequence just by looking at natural language, any more than you can find natural language semantics by looking at your favorite logic. Logic and natural language are autonomous. All the same, if you already have some prior, perhaps rough or general, grip on logic, you can find lots of interesting facts in natural language, which you might use in further theorizing. This could even lead you to extract a full-blown logical consequence relation from a natural language, but you have to start with some logic if you want to get some logic.

References

Abney, S., 1987. *The English Noun Phrase in its Sentential Aspect*. Ph.D. dissertation, MIT.

Anderson, A. R. and N. D. Belnap, 1975. *Entailment: The Logic of Relevance and Necessity*, vol. I. Princeton: Princeton University Press.

Anderson, A. R., N. D. Belnap, and J. M. Dunn, 1992. *Entailment: The Logic of Relevance and Necessity*, vol. II. Princeton: Princeton University Press.

Anderson, S. R., 1971. On the role of deep structure in semantic interpretation. *Foundations of Language* **7**:387–396.

Antonelli, G. A., 2005. *Grounded Consequence for Defeasible Logic*. Cambridge: Cambridge University Press.

Aoun, J. and Y.-h. A. Li, 1993. *Syntax of Scope*. Cambridge: MIT Press.

Asher, N. and A. Lascarides, 2003. *Logics of Conversation*. Cambridge: Cambridge University Press.

Barker, C., 2002. Continuations and the nature of quantification. *Natural Language Semantics* **10**:211–242.

Barwise, J. and R. Cooper, 1981. Generalized quantifiers and natural language. *Linguistics and Philosophy* **4**:159–219.

Barwise, J. and S. Feferman (eds.), 1985. *Model-Theoretic Logics*. Berlin: Springer.

Beall, J. C. and B. C. van Fraassen, 2003. *Possibilities and Paradox*. Oxford: Oxford University Press.

Beall, J. C. and G. Restall, 2006. *Logical Pluralism*. Oxford: Oxford University Press.

———, 2009. Logical consequence. In E. N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*. Fall 2009 edn.

Beghelli, F. and T. Stowell, 1997. Distributivity and negation: The syntax of *Each* and *Every*. In A. Szabolcsi (ed.), *Ways of Scope Taking*, pp. 71–107. Dordrecht: Kluwer.

van Benthem, J., 1986. *Essays in Logical Semantics*. Dordrecht: Reidel.

———, 1991. *Language in Action*. Amsterdam: North-Holland.

Bonnay, D., 2008. Logicality and invariance. *Bulletin of Symbolic Logic* **14**:29–68.

Carnap, R., 1947. *Meaning and Necessity*. Chicago: University of Chicago Press.

Carpenter, B., 1997. *Type-Logical Semantics*. Cambridge: MIT Press.

Carroll, L., 1960. *The Annotated Alice*. Meridian. Introduction and Notes by Martin Gardner.

Casari, E., 1987. Comparative logics. *Synthese* **73**:421–449.

Chierchia, G. and S. McConnell-Ginet, 2000. *Meaning and Grammar: An Introduction to Semantics*. Cambridge: MIT Press, 2nd edn.

Church, A., 1940. A formulation of the simple theory of types. *Journal of Symbolic Logic* **5**:56–68.

Cresswell, M. J., 1973. *Logics and Languages*. London: Methuen.

———, 1978. Semantics and logic. *Theoretical Linguistics* **5**:19–30.

Davidson, D., 1967. Truth and meaning. *Synthese* **17**:304–323. Reprinted in Davidson (1984).

———, 1984. *Inquiries into Truth and Interpretation*. Oxford: Oxford University Press.

Dowty, D. R., 1991. Thematic proto-roles and argument selection. *Language* **67**:547–619.

Dowty, D. R., R. E. Wall, and S. Peters, 1981. *Introduction to Montague Semantics*. Dordrecht: Reidel.

Dummett, M., 1991. *The Logical Basis of Metaphysics*. Cambridge: Harvard University Press.

Etchemendy, J., 1988. Tarski on truth and logical consequence. *Journal of Symbolic Logic* **53**:51–79.

———, 1990. *The Concept of Logical Consequence*. Cambridge: Harvard University Press.

Evans, G., 1976. Semantic structure and logical form. In G. Evans and J. McDowell (eds.), *Truth and Meaning*, pp. 199–222. Oxford: Oxford University Press.

Feferman, S., 1999. Logic, logics, and logicism. *Notre Dame Journal of Formal Logic* **40**:31–54.

Fillmore, C. J., 1968. The case for case. In E. Bach and R. T. Harms (eds.), *Universals in Linguistic Theory*, pp. 1–88. New York: Holt, Rinehart and Winston.

Fodor, J. D. and I. A. Sag, 1982. Referential and quantificational indefinites. *Linguistics and Philosophy* **5**:355–398.

Frege, G., 1879. *Begriffsschrift, eine der arithmetischen nachgebildete Formelsprache des reinen Denkens*. Halle: Nebert. References are to the translation as "Begriffsschrift, a Formal Language, Modeled upon that of Arithmetic, for Pure Thought" by S. Bauer-Mengelberg in van Heijenoort (1967).

———, 1891. *Function und Begriff*. Jena: Pohle. References are to the translation as "Function and Concept" by P. Geach in Frege (1984).

———, 1984. *Collected Papers on Mathematics, Logic, and Philosophy*. Oxford: Basil Blackwell. Edited by B. McGuinness.

Frigg, R. and S. Hartmann, 2009. Models in science. In E. N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*. Summer 2009 edn.

Fukui, N., 1995. *Theory of Projection in Syntax*. Stanford: CSLI Publications.

Gamut, L. T. F., 1991. *Logic, Language, and Meaning*, vol. 2: Intensional Logic and Grammar. Chicago: University of Chicago Press. 'Gamut' is a pseudonym for van Benthem, Groenendijk, de Jongh, Stokhof, and Verkuyl.

Gentzen, G., 1935. Untersuchungen über das logische Schliessen. *Mathematische Zeitschrift* **39**:176–210. References are to the translation as "Investigations into logical deduction" in Szabo (1969).

Gödel, K., 1930. Die Vollständigkeit der Axiome des logischen Funktionenkalküls. *Monatshefte für Mathematik und Physik* **37**:349–360. References are to the translation as "The completeness of the axioms of the functional calculus of logic" by S. Bauer-Mengelberg in Gödel (1986).

———, 1986. *Collected Works*, vol. I. Oxford: Oxford University Press. Edited by S. Feferman, J. W. Dawson Jr., S. C. Kleene, G. H. Moore, R. M. Solovay, and J. van Heijenoort.

Grice, P., 1975. Logic and conversation. In P. Cole and J. L. Morgan (eds.), *Speech Acts*, vol. 3 of *Syntax and Semantics*, pp. 41–58. New York: Academic Press. Reprinted in Grice (1989).

———, 1989. *Studies in the Way of Words*. Cambridge: Harvard University Press.

Grimshaw, J., 2005. Extended projection. In *Words and Structure*, pp. 1–73. Stanford: CSLI Publications.

Hale, K. and S. J. Keyser, 1987. A view from the middle. Lexicon Project Working Papers 10, Center for Cognitive Science, MIT.

van Heijenoort, J. (ed.), 1967. *From Frege to Gödel: A Source Book in Mathematical Logic, 1879–1931*. Cambridge: Harvard University Press.

Heim, I. and A. Kratzer, 1998. *Semantics in Generative Grammar*. Oxford: Blackwell.

Higginbotham, J., 1986. Linguistic theory and Davidson's program in semantics. In E. Lepore (ed.), *Truth and Interpretation: Perspectives on the Philosophy of Donald Davidson*, pp. 29–48. Oxford: Blackwell.

———, 1988. Contexts, models, and meanings: A note on the data of semantics. In R. Kempson (ed.), *Mental Representations: The Interface Between Language and Reality*, pp. 29–48. Cambridge: Cambridge University Press.

———, 1989a. Elucidations of meaning. *Linguistics and Philosophy* **12**:465–517.

———, 1989b. Knowledge of reference. In A. George (ed.), *Reflections on Chomsky*, pp. 153–174. Oxford: Basil Blackwell.

Higginbotham, J. and R. May, 1981. Questions, quantifiers and crossing. *Linguistic Review* **1**:41–79.

Hirschberg, J., 1985. *A Theory of Scalar Implicature*. Ph.D. dissertation, University of Pennsylvania.

Horty, J. F., 2001. Nonmonotonic logic. In L. Goble (ed.), *The Blackwell Guide to Philosophical Logic*, pp. 336–361. Oxford: Blackwell.

Kadmon, N., 2001. *Formal Pragmatics*. Oxford: Blackwell.

Keenan, E. L. and J. Stavi, 1986. A semantic characterization of natural language determiners. *Linguistics and Philosophy* **9**:253–326. Versions of this paper were circulated in the early 1980s.

Kennedy, C. and L. McNally, 2005. Scale structure, degree modification, and the semantics of gradable predicates. *Language* **81**:345–381.

Klein, E. and I. A. Sag, 1985. Type-driven translation. *Linguistics and Philosophy* **8**:163–201.

Kneale, W. and M. Kneale, 1962. *The Development of Logic*. Oxford: Clarendon.

Kretzmann, N., A. Kenny, and J. Pinborg (eds.), 1982. *The Cambridge History of Later Medieval Philosophy*. Cambridge: Cambridge University Press.

Kripke, S. A., 1963. Semantical considerations on modal logic. *Acta Philosophica Fennica* **16**:83–94.

Lambek, J., 1958. The mathematics of sentence structure. *American Mathematical Monthly* **65**:154–170.

Larson, R. and G. Segal, 1995. *Knowledge of Meaning*. Cambridge: MIT Press.

Lepore, E., 1983. What model-theoretic semantics cannot do. *Synthese* **54**:167–187.

Levin, B., 1993. *English Verb Classes and Alternations*. Chicago: University of Chicago Press.

Levin, B. and M. Rappaport Hovav, 1995. *Unaccusativity: At the Syntax-Lexical Semantics Interface*. Cambridge: MIT Press.

———, 2005. *Argument Realization*. Cambridge: Cambridge University Press.

Lewis, C. I. and C. H. Langford, 1932. *Symbolic Logic*. New York: Century.

Lewis, D., 1970. General semantics. *Synthese* **22**:18–67. Reprinted in Lewis (1983).

———, 1983. *Philosophical Papers*, vol. 1. Oxford: Oxford University Press.

Lindström, P., 1966. First order predicate logic with generalized quantifiers. *Theoria* **32**:186–195.

Lycan, W. G., 1989. Logical constants and the glory of truth-conditional semantics. *Notre Dame Journal of Formal Logic* **30**:390–400.

MacFarlane, J., 2009. Logical constants. In E. N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*. Fall 2009 edn.

Mancosu, P., 2010. Fixed- versus variable-domain interpretations of Tarski's account of logical consequence. *Philosophy Compass* **5**:745–759.

Mautner, F. I., 1946. An extension of Klein's erlanger program: Logic as invariant-theory. *American Journal of Mathematics* **68**:345–384.

Milsark, G., 1974. *Existential Sentences in English*. Ph.D. dissertation, MIT. Published by Garland, New York, 1979.

Montague, R., 1968. Pragmatics. In R. Klibansky (ed.), *Contemporary Philosophy: A Survey*, pp. 102–122. Florence: La Nuova Italia Editrice. Reprinted in Montague (1974).

———, 1970. Universal grammar. *Theoria* **36**:373–398. Reprinted in Montague (1974).

———, 1973. The proper treatment of quantification in ordinary English. In J. Hintikka, J. Moravcsik, and P. Suppes (eds.), *Approaches to Natural Language*, pp. 221–242. Dordrecht: Reidel. Reprinted in Montague (1974).

———, 1974. *Formal Philosophy: Selected Papers of Richard Montague*. New Haven: Yale University Press. Edited by R. H. Thomason.

Moortgat, M., 1997. Categorial type logics. In J. van Benthem and A. ter Meulen (eds.), *Handbook of Logic and Language*, pp. 93–177. Amsterdam: Elsevier.

Mostowski, A., 1957. On a generalization of quantifiers. *Fundamenta Mathematicae* **44**:12–36.

Partee, B. H., 1979. Semantics—mathematics or psychology? In A. Bäuerle, U. Egli, and A. von Stechow (eds.), *Semantics from Different Points of View*, pp. 1–14. Berlin: Springer.

Peters, S. and D. Westerståhl, 2006. *Quantifiers in Language and Logic*. Oxford: Clarendon.

Pietroski, P. M., 2005. *Events and Semantic Architecture*. Oxford: Oxford University Press.

Pinker, S., 1989. *Learnability and Cognition*. Cambridge: MIT Press.

Portner, P. and B. H. Partee (eds.), 2002. *Formal Semantics: The Essential Readings*. Oxford: Blackwell.

Prawitz, D., 1974. On the idea of a general proof theory. *Synthese* **27**:63–77.

———, 2005. Logical consequence from a constructivist point of view. In S. Shapiro (ed.), *The Oxford Handbook of Philosophy of Mathematics and Logic*, pp. 671–695. Oxford: Oxford University Press.

Priest, G., 2008. *An Introduction to Non-Classical Logic*. Cambridge: Cambridge University Press, 2nd edn.

Quine, W. V. O., 1958. Speaking of objects. *Proceedings and Addresses of the American Philosophical Association* **31**:5–22. Reprinted in Quine (1969).

———, 1959. *Methods of Logic*. New York: Holt, revised edn.

———, 1960. *Word and Object*. Cambridge: MIT Press.

———, 1969. *Ontological Relativity and Other Essays*. New York: Columbia University Press.

———, 1986. *Philosophy of Logic*. Cambridge: Harvard University Press, 2nd edn.

Read, S., in press. The medieval theory of consequence. *Synthese* .

Restall, G., 2000. *An Introduction to Substructural Logics*. London: Routledge.

Russell, B., 1903. *Principles of Mathematics*. Cambridge: Cambridge University Press.

Sag, I. A., T. Wasow, and E. M. Bender, 2003. *Syntactic Theory: A Formal Introduction*. Stanford: CSLI Publications, 2nd edn.

Scott, D., 1970. Advice on modal logic. In K. Lambert (ed.), *Philosophical Problems in Logic*, pp. 143–173. Dordrecht: Reidel.

Shapiro, S., 1991. *Foundations without Foundationalism: A Case for Second-Order Logic*. Oxford: Clarendon.

Sher, G., 1991. *The Bounds of Logic*. Cambridge: MIT Press.

Skolem, T., 1922. Einige Bemerkungen zur axiomatischen Begründung der Mengenlehre. In *Matematikerkongressen i Helsingfors den 4–7 Juli 1922, Den femte skandinaviska matematikerkongressen, Redogörelse*, pp. 217–232. Helsinki: Akademiska Bokhandeln. References are to the translation as "Some remarks on axiomatized set theory" by S. Bauer-Mengelberg in van Heijenoort (1967).

Smith, R., 1995. Logic. In J. Barnes (ed.), *The Cambridge Companion to Aristotle*, pp. 27–65. Cambridge: Cambridge University Press.

Speas, M., 1990. *Phrase Structure in Natural Language*. Dordrecht: Kluwer.

Strawson, P. F., 1950. On referring. *Mind* **59**:320–344. Reprinted in Strawson (1971).

———, 1971. *Logico-Linguistic Papers*. London: Methuen.

Szabo, M. E. (ed.), 1969. *The Collected Papers of Gerhard Gentzen*. Amsterdam: North-Holland.

Tarski, A., 1935. Der Wahrheitsbegriff in den formalisierten Sprachen. *Studia Philosophica* **1**:261–405. References are to the translation as "The concept of truth in formalized languages" by J. H. Woodger in Tarski (1983). Revised version of the Polish original published in 1933.

———, 1936. O pojciu wynikania logicznego. *Przegląd Filozoficzny* **39**:58–68. References are to the translation as "On the concept of logical consequence" by J. H. Woodger in Tarski (1983).

———, 1983. *Logic, Semantics, Metamathematics*. Indianapolis: Hackett, 2nd edn. Edited by J. Corcoran with translations by J. H. Woodger.

———, 1986. What are logical notions? Text of a 1966 lecture edited by J. Corcoran. *History and Philosophy of Logic* **7**:143–154.

Thomason, R. H., 1974. Introduction. In R. H. Thomason (ed.), *Formal Philosophy: Selected Papers of Richard Montague*, pp. 1–69. New Haven: Yale University Press.

Westerståhl, D., 1985. Logical constants in quantifier languages. *Linguistics and Philosophy* **8**:387–413.

———, 1989. Quantifiers in formal and natural languages. In D. Gabbay and F. Guenthner (eds.), *Handbook of Philosophical Logic*, vol. IV, pp. 1–131. Dordrecht: Kluwer.

Whitehead, A. N. and B. Russell, 1925–27. *Principia Mathematica*. Cambridge: Cambridge University Press, 2nd edn.

Zimmermann, T. E., 1999. Meaning postulates and the model-theoretic approach to natural language semantics. *Linguistics and Philosophy* **22**:529–561.