



Reconciling NSM and Formal Semantics

Avery D. Andrews

To cite this article: Avery D. Andrews (2015): Reconciling NSM and Formal Semantics, Australian Journal of Linguistics, DOI: [10.1080/07268602.2016.1109431](https://doi.org/10.1080/07268602.2016.1109431)

To link to this article: <http://dx.doi.org/10.1080/07268602.2016.1109431>



Published online: 15 Dec 2015.



Submit your article to this journal [↗](#)



Article views: 1



View related articles [↗](#)



View Crossmark data [↗](#)

Reconciling NSM and Formal Semantics*

AVERY D. ANDREWS

Australian National University

(Accepted 23 August 2015)

Formal semantics and Natural Semantic Metalanguage are widely held to be radically incompatible as ways to study meaning in natural language. Here I will show that they can to some extent be reconciled. In particular, for linguists working with NSM, formal semantics can be viewed as providing mathematical accounts of some of the same phenomena that NSM studies, such as entailment, and for the formal semanticist, NSM offers a small target in the form of mini-languages that exhibit the essential logico-semantic features of full natural languages, such as extensionality, intensionality and hyperintensionality, and algebraic principles such as transitivity, symmetry etc. or various of the primes. Therefore, although these two approaches are likely to remain distinct enterprises for the foreseeable future, some intercommunication is possible and indeed desirable.

Keywords: Natural Semantic Metalanguage; NSM; Formal Semantics; Model Theoretic Semantics; Inference Rules; Algebraic Semantics

1. Overview

The approach to semantics known as Natural Semantic Metalanguage (NSM), originated by Anna Wierzbicka in the early 1970s (Wierzbicka 1972), and formal semantics, as pioneered by Richard Montague a little bit earlier (Montague 1970a, 1970b), are usually seen as mutually exclusive research programmes. Here, I will argue that they can and should be regarded as complementary to at least some extent, although this calls for reconsideration of some assumptions often presented as foundational. There are at least two major issues of this kind, one being the role of mathematical

*I would like to acknowledge many extremely useful comments from three anonymous (*Australian Journal of Linguistics*) reviewers, in addition to comments on earlier versions from Maïa Ponsonnet, Zhengdao Ye and Andras Kornai, and also especially Barbara Partee for many discussions over the years, and comments on what might be regarded as preliminaries to this paper. But they are not responsible for my errors.

formulation, the other the adoption of a ‘realist’ as opposed to ‘antirealist’ view of meaning. Mathematical formulation is fundamental to formal semantics, but traditionally avoided in NSM. So one of the goals of this paper is to show how one can begin to apply some of the mathematical methods of formal semantics to NSM (emphasis on ‘begin’). This enterprise should be of some interest to formal semanticists for at least two reasons, one ‘good news’, the other, ‘bad news’. The bad news reason is that since NSMs are subsets of natural languages, with restricted vocabulary (the ‘primes’) and syntax, it follows that if it is impossible to do formal semantics for a proposed NSM for a language, it is impossible to do formal semantics at all for that language (and, very likely, any natural language, to the extent that the universality claims of NSM work out). The good news reason, on the other hand, hinges on the small size combined with universality of NSM: because NSMs are small, they are far more feasible as targets for formalization than entire languages, and because they are supposed to be constructed out of essentially universal elements, with serious investigation of their universality, formalizing one of them would be a big step towards formalizing all of them, strengthening the case that natural language semantics is in general capable of formalization.

Why linguists working in NSM should be interested in formalization is less easy to explain (in terms of their point of view, which is rather strange to people steeped in late-twentieth century cognitivism), but I suggest the following: a convincing demonstration that natural language semantics either can or cannot be formalized would be a very important result of great cultural significance, indicating for example whether or not the project of creating human-like artificial intelligence is possible using current mathematical ideas and engineering techniques. Especially interesting would be a solidly established rather than anecdotally advocated case for the negative conclusion that NL semantics cannot be formalized. But one cannot demonstrate something of this nature with programmatic announcements and the exhibition of a few sample difficulties. The failure of serious attempts with apparently reasonable prospects for success is required. The small size of NSMs makes them ideal laboratories for investigating this question.

The other major issue, realism vs. antirealism, is rhetorically prominent, but, I claim, essentially a distraction that can and should be set aside, at least for the present. Semantic realism (i.e. externalism) claims that the central task of semantics is to connect language to the world, especially in terms of our ability to assess the truth of statements (see Lewis (1972) and Cresswell (1978) for classic expositions). Antirealism (internalism) on the other hand claims that this project is impossible and senseless as a goal for semantics, which should instead focus on intuitively perceivable properties of and relations between sentences, such as anomaly and entailment (somewhat comparable to ‘grammaticality’ in syntax, although with significant differences in detail). This is in practice the standard position of NSM, and in the generative world, received a comprehensive and influential formulation in Chapter 1 of Katz (1972), and a rarely answered philosophical defence in LePore (1983). Ray Jackendoff is perhaps the most prominent generative advocate of this position today; see

Jackendoff (1985) for a critique of realist semantics (specifically, the realist commitments of Barwise and Perry's (1983) Situation Theory), and Barwise and Perry (1985), especially pages 112–137, for Barwise and Perry's evidently baffled and rather frustrated reply.

We are obviously not going to settle this here, but can note that in some sense, realism is at least partially correct: it seems very unlikely that language could have developed as it has if it lacked real-world significance. On the other hand, linguists are not very well equipped to study this with their native methods; what they can investigate with these are patterns of entailment and similar relationships between sentences, as proposed by Katz, and called by Larson and Segal (1995) 'Logico-Semantic Properties and Relations' (Katz seems to have provided no generic name for them). I will call them 'Meaning-Based Properties and Relations' (MBPRs) in an attempt to avoid complex technical terms when possible. Of the MBPRs, entailment and the closely related consistency and contradiction (all truth-based) are the most closely studied, with many others, such as presupposition and 'appropriate answer to a question' getting significant attention.¹

The issue of realism is highly relevant to the status of model theory, the premier tool of formal semantics, which is frequently justified by its claimed capacity to explain the relationship of language to the world. The position here will be that whether or not it can actually do this, model theory is (a) an excellent mathematical tool for characterizing many MBPRs, especially entailment (although in a non-constructive manner); and (b) a highly successful 'context of inspiration', whereby informal reflection on aspects of everyday language use can provide motivation for features of mathematical theories of entailment, as well as the occasional enlightening experiment. We don't need to fully understand why this is the case in order to benefit from it.

In the first section, I will discuss formal semantics, starting with model theory, and moving on to deduction, algebraic semantics and 'support'. A major aim will be to try to make model theoretic semantics look worthwhile to people who reject the case for realism. Since NSM researchers have by and large already rejected the rather cursory presentations of the justifications for model theoretic semantics that are found in textbooks, we take a somewhat extended and perhaps idiosyncratic tour through foundational issues. Its goal is not so much to persuade as to encourage further thought.

Then, in the second section, I will present a brief sketch of certain aspects of NSM, including especially comparisons of the earlier versions of the theory (up to the mid-1980s) with the later and current ones. The aim is not advocacy, but some basic explanation.

Finally in the third section I will show how, in spite of the differences between the approaches, certain aspects of formal semantics, especially model theory, can be applied to NSM, considering first the rather surprising phenomena of truth-functionality, extensionality and intensionality (as opposed to 'hyperintensionality', which I

¹ Katz himself had little to say about entailment as such, focusing on 'analyticity', without saying anything that people appear to have managed to understand as to how this was supposed to be different from entailment.

suggest is in fact the default expectation), and then surveying a range of entailment laws that seem to be obeyed by various NSM ‘primes’ (hypothesized fundamental, universal concepts behind human languages, as explained in Section 2). This is only a rather small start on linking up these two research programmes, but it is enough to reveal that it is not a completely impossible thing to work on.

2. Formal Semantics

We begin with a brief and incomplete sketch of model theoretic semantics (MTS), whose purpose is to indicate how it provides a mathematical theory of MBPRs independently of any claims about realism in semantics. Then we consider the deductive alternative, and discuss some possible reasons why it has appeared to ‘fail to thrive’ in spite of being in many ways more consistent with the usual assumptions of generative grammar (and, also, NSM). Then, after returning to MTS to discuss its benefits in comparison to deductive approaches, we consider the further topics of algebraic semantics, and a notion of ‘support’, which can be regarded as an attenuated version of realism. The first two subsections can be skipped or skimmed by readers with a basic knowledge of formal semantics. Much of the content of the later subsections will also be familiar to such readers, but is intended to indicate what kinds of adjustments might be useful for better integration of formal semantics with NSM.

2.1. Model Theory for Antirealists

A simple MTS for a fragment of a natural language can be constructed like this. We assume that our formal fragment contains n -place predicates for various integers n , together with some ‘logical words’ such as *and*, *not* and *if* (and, for a more advanced fragment than we will attempt here, *some* and/or *every*), and some proper names such as *John* and *Mary*. An ‘interpretation’ I for the language is then constituted by the following:

- (1) a. A set A to serve as a ‘universe of discourse’ (collection of entities).
- b. A member of A to serve as ‘extension’ of each proper name (in more advanced systems with quantifiers, variables will also range over the members of A).
- c. For each n -ary predicate, an n -ary relation, that is, a set of n -tuples on A , to serve its extension.

So, if our universe of discourse is $\{j, b, m, s\}$, and the extension of the two-place predicate *Love* is $\{\langle m, j \rangle, \langle j, m \rangle\}$, we can take this as representing a situation where two people love each other, but a third is not involved.

We now define ‘truth relative to an interpretation’ inductively, by extending (1) to an assignment of truth-values to the formulas constructed from the basic symbols of

the language, together with further rules for the logical words. For the n -place predicates, we have:

- (2) If P is an n -place predicate, and $F = P(t_1, \dots, t_n)$ is a formula, then $I(F) = 1$ if $\langle I(t_1), \dots, I(t_n) \rangle \in I(P)$, otherwise $I(F) = 0$.

For example, if $I(\text{John}) = j$, $I(\text{Mary}) = m$ and $I(\text{Love}) = \{\langle j, m \rangle\}$, then $I(\text{Love}(j, m)) = 1$, $I(\text{Love}(m, j)) = 0$ (romantic purgatory).

According to this technical account, for 1-place predicates such as *Boy* and *Girl*, the interpretations will be sets of 1-tuples of elements of A , e.g. $I(\text{Boy}) = \{\langle j \rangle\}$, $I(\text{Girl}) = \{\langle m \rangle, \langle s \rangle\}$, which are effectively equivalent to subsets of A , as usually presented in introductory logic textbooks.

For the logical words *and* and *not*, we assume syntax rules that combine *not* with one formula and *and* with two, putting parentheses around the result in the latter case to forestall ambiguity, with these rules:

- (3) a. If $F = \text{not } G$ is a formula, then $I(F) = 1$ if $I(G) = 0$, $I(F) = 0$ otherwise.
 b. If $F = (G \text{ and } H)$ is a formula, then $I(F) = 1$ if $I(G) = 1$ and $I(H) = 1$, otherwise $I(F) = 0$.

The standard logical word *or* on the other hand, is not accepted as a prime (and can be either defined or explicated in various ways). The case of *if* is more complex, and will be deferred.

These rules work by climbing up the syntax tree, assigning values to larger constituents on the basis of those of the smaller ones, starting with the n -place predicates. So for example, what values do these sentences get under the ‘romantic purgatory’ interpretation above?

- (4) a. $\text{Love}(\text{Mary}, \text{John}) \text{ and } \text{Love}(\text{John}, \text{Mary})$
 b. $(\text{not } \text{Love}(\text{Mary}, \text{John}) \text{ and } \text{Love}(\text{John}, \text{Mary}))$
 c. $\text{not } (\text{Love}(\text{Mary}, \text{John}) \text{ and } \text{Love}(\text{John}, \text{Mary}))$

In spite of its incompleteness (no account of quantifiers being why it falls short of constituting first order logic), this example demonstrates important features of how MTS works, but also a fundamental problem for its justification in much of the formal semantics literature.

MTS is supposed to help explain truth by showing how the truth of sentences is determined by how the world is, but, arguably, as discussed by LePore (1983), this is not happening here. Rather, we are connecting truth and falsity of sentences to certain mathematical structures. To connect the sentences to the world, we would have to move on to connect the mathematical structures to the world, and there does not appear to be any well worked-out and generally accepted story about how

to do this (although below I will introduce a limited version of this that I will call ‘support’).

So, one can say, however persuasive the philosophical position of realist semantics might be, model theoretic semantics of natural language does not fully implement it. That is, it is in the end unexplained by formal semantics how a sentence such as *Everybody went to Mooseheads* is false if the people at the philosophy seminar actually all went to The Wig and Pen. This might appear at first to be a fatal objection, but if we pursue things a bit further we find that it isn’t, because many useful things survive it.

This is because what is typically done with a model theory, once it is set up, is to give a characterization of at least entailment, and perhaps some of the other MBPRs. A classic example of entailment would be the following syllogism where Σ , the ‘premises’ of the syllogism, are above the horizontal line, S , its ‘conclusion’, below it:

- $$(5) \quad \begin{array}{l} \text{Socrates is a man} \\ \text{All men are mortal} \\ \hline \therefore \text{Socrates is mortal} \end{array}$$

Generalizing, we can define entailment conceptually like this:

- (6) A set of sentences Σ entails a sentence S if any person who accepts the sentences of Σ as true ‘must’ also accept that S is true.

The exact interpretation of *must* here is something that one can debate, but I contend that it cashes out to a conviction that if somebody accepts the sentences of Σ but rejects S , there is absolutely no point in continuing a discussion with them. They can be written off as hopeless conversational partners. Note furthermore that one can recognize specific instances of entailment without having any kind of theory of how they arise (which is exactly our position with example (5) above, since we have not suggested any theory of *all*).

In MTS, one provides a mathematical characterization of entailment as follows:

- (7) A set of sentences Σ entails a sentence S if every interpretation that makes all the sentences of Σ true also makes S true.

Here, (7) is called a ‘characterization’ rather than a ‘definition’ because, while the definition as given in (6) is supposed to tell us what kinds of things we ought to accept as proposed mathematical characterizations of entailment, (7) presents a general format for producing such characterizations (there is another technique, inference rules, as we will discuss soon).

An essential point is that the truth or falsity with respect to any specific interpretation becomes irrelevant here; what matters is the pattern of truth-assignments across all of the sentences of the language. The problem of connecting any particular

assignment to the details of how things are in the real world then becomes something which we don't have to concern ourselves with (but can, if we find an interesting aspect of it that is accessible to our research methods).

That MTS provides a way of characterizing one MBPR does not of course mean that it is the best way of characterizing that MBPR or any of the others. To judge its relative merits, we have to look at the alternatives, which, I contend, means deductive a.k.a. proof-theoretical formulations (Davidsonian 'truth conditional' semantics as developed in detail in Larson and Segal (1995) might be seen as an alternative, but Larson and Segal in fact endorse the use of a deductive account of the MBPRs, but don't make any specific proposal). Nevertheless, it does show that we can use MTS while being agnostic about or hostile to realism; a substantial piece of MTS-based work that takes an explicitly agnostic point of view is Keenan and Faltz (1985).

2.2. The Deductive Alternative

In linguistics as well as in logic, where the study of these questions originated, the original accounts of what I am calling MBPRs were not formulated with model theory, but rather with what can be reasonably called principles of deduction, whereby the relationships between sentences are described in terms of their forms. For example, glossing over many details, a basic start on syllogisms can be made by proposing that any argument of the following form is 'valid', that is, if we accept the premises above the line, we 'must' accept the conclusion below it:

$$(8) \quad \begin{array}{l} \text{Every } P \text{ is } Q \\ \hline e \text{ is } P \\ \hline \therefore e \text{ is } Q. \end{array}$$

Introductory logic textbooks present deductive systems that you can do useful things with, and Kneale and Kneale (1962) is the standard source for the development of the subject from classical antiquity to the mid-twentieth century, shortly before the beginnings of formal semantics.

In contemporary linguistics, this general form of approach (with completely different and highly unsatisfactory details) was introduced by Katz and Fodor (1964), in their attempt to characterize 'analyticity' in terms of a subset relation between elements in the semantic representation of sentences, and related notions such as 'contradictoriness' along similar lines. Amongst the deficiencies of this formulation was that it had no account of *n*-place predicates and their arguments, and therefore was incapable of distinguishing the meanings of *John loves Mary* and *Mary loves John*. Later revisions such as those in Katz (1966) and Katz (1972) remedy this deficiency and were considerably better, but perhaps in part because of the bad impression made by the original version (Partee 2004), they never attained a significant following. Nevertheless, Katz (1972: ch. 1) stands out as a very clear formulation of the idea that a major goal of semantics should be accounting for the MBPRs.

Although this goal has been widely accepted, at least implicitly,² neither Katz's technical execution nor any of the others that have appeared in 'representational semantics' have gone very far. Jackendoff's last inference rule was for example proposed in Jackendoff (1983), and similar approaches in linguistics do not appear to have done any better, at least prior to the rise of recent investigations of 'Natural Logic',³ which has significantly different goals than standard formal semantics. Lakoff's (1970) proposal to use standard logical deduction rules applying to generative-semantics' 'deep structures' likewise does not appear to have led to substantial results.

The failure to thrive of all of these approaches is something of a puzzle, because not only were they on the scene prior to MTS-based formal semantics, but they are also more in accord with the basic ideas of generative grammar: a deductive account of entailment does, after all, constitute a theory of how it is in principle possible for a native speaker of a language to have an entailment intuition, namely, by constructing the deduction and construing this as a basis for accepting the entailment. A model theoretic account on the other hand can't do this, because the definition involves *all possible interpretations* (which will in general be infinite in number), which the language user obviously cannot run through and check individually.

Rather, an MTS account of entailment constitutes what in mathematics would be called a 'non-constructive' characterization of entailment, a characterization that is in some sense logically or intuitively valid, but doesn't tell you anything directly about how to recognize or construct the things that it characterizes (but whose details can often, but do not always, provide hints for the construction of such techniques).

Nevertheless, in linguistics, MTS started developing rapidly, and has continued to do so for decades. My suggestion is that the basic reason for this is that proof theory (the currently standard name for the mathematical investigation of deductive systems) is considerably harder than model theory, at least with respect to linguistic applications, and that this difference in difficulty was more extreme in the early days of formal semantics, when proof theory was less developed, and there were far fewer resources available from which to learn anything about it. Linguists trying to make up deductive systems on their own, in particular, were I think almost certain not to get very far. To develop this point, we need to examine in closer detail how such accounts work. We will do this with conventional logical rules of deduction, since these are very well understood.

Deductive accounts of an MBPR such as, for example, entailment, work by formulating the property or relation in question as a relation over syntactic structures. In classic symbolic logic, these structures were 'formulas' in invented artificial languages, with no formalized connection to natural language sentences, while in formal semantics, they are linguistic structures that are abstract to at least some extent (involving, at

² Explicit acknowledgements seem to be rare, absent for example from Chierchia and McConnell-Ginet (2000: ch. 1), although many of the same MBPRs are discussed in a similar way.

³ As presented for example in <http://nlp.stanford.edu/projects/natlog.shtml> (accessed 6 June 2015).

a minimum, some account of constituent structure). In this presentation we will not worry about the details of the formulas/abstract structures.

In contemporary work, a popular notation for expressing entailment relations is the ‘natural deduction tree format’, where the premises are written above a line, and the/a conclusion below. From premises consisting of two clauses combined by *and*, for example, we can derive either of the conjuncts as a conclusion; such a rule is called an ‘elimination’ rule in Natural Deduction:

$$(9) \quad \frac{A \text{ and } B}{A} \text{and-elim} \qquad \frac{A \text{ and } B}{B} \text{and-elim}$$

Or, given *A* and *B* as premises, we can produce their conjunction as conclusion, called *and*-introduction:

$$(10) \quad \frac{A \quad B}{A \text{ and } B} \text{and-intr}$$

Such rules can be applied in combinations to produce additional results; for example we can derive *B and A* from *A and B* like this:

$$(11) \quad \frac{\frac{A \text{ and } B}{B} \text{and-elim} \qquad \frac{A \text{ and } B}{A} \text{and-elim}}{B \text{ and } A} \text{and-intr}$$

The possible results get more interesting as we add more rules, so that for example if suitable rules and representations for quantifiers and their associated ‘variables’ are adopted, we can derive results such as the following three-premise argument, which the Ancient Greeks could not accommodate:

$$(12) \quad \begin{array}{l} \text{Every boy loves every girl} \\ \text{John is a boy} \\ \text{Mary is a girl} \\ \therefore \text{John loves Mary} \end{array}$$

So deduction is certainly in general a workable approach to characterizing MBPRs, but suffers from at least two substantial problems (especially for linguists several decades ago), one of which is the requirement for very large numbers of arbitrary decisions, and the other a discouragingly high level of mathematical difficulty.

2.3. Problems with Deduction

The arbitrariness problem consists of the fact that there are a very large number of ways in which deductive systems can be set up. Anybody with a reasonable amount

of basic knowledge about logic will be aware of at least the following different major techniques:

- (13) a. Hilbert Systems (the oldest, and most still most powerful, but not very intuitive).
 b. Natural Deduction, in any of a large number of variant formulations (usefully surveyed by Pelletier (1999); the main alternative to the ‘tree-style’ format used above, common in logic textbooks, is ‘Fitch-style’; there is also ‘sequent style’ natural deduction).
 c. Gentzen Sequents (different from sequent-style natural deduction, although both were invented by Gentzen).
 d. Beth Tableaux.

There are presumably additional options known to specialists, as well as sub-variants of the above systems. Nowadays, a great deal is known about the equivalences and other relationships between these systems, but it still takes some time to learn a reasonable amount of this material, and it would have taken much longer in the 1970s, when formal semantics was taking shape, because there was much less in the way of accessible textbooks and tutorial papers, and less was known overall. The other problem, difficulty, compounds with arbitrariness. In order to have empirical content, a theory of entailment has to characterize some collections of premises as entailing some possible conclusions, others as not. In a deductive approach, the former is relatively easy: to show that some premises entail a conclusion, produce a derivation of the former from the latter (actually, in fact, what is easy is to check a putative fulfilment of this goal; finding a proof can in fact be rather hard). But the latter requires showing that there is *no* proof of the conclusion from the premises, which can be difficult (especially for linguists without a lot of mathematical aptitude and training).

This problem is exacerbated by the fact that linguists can’t confine themselves to the somewhat limited range of deductive patterns employed in mathematical discussion, but must consider a much wider range of naturally occurring ones, such as that *John and Bill each carried a piano downstairs* entails both *John carried a piano downstairs*, and *Bill carried a piano downstairs* but *John and Bill carried a piano downstairs (together)* does not. How can you be sure that the new deductive rule you’re contemplating won’t cause your system to melt down in contradictions? MTS seems on the basis of experience to ameliorate the problem of arbitrariness, and to provide a real solution to the problem of accidental introduction of contradictions, since it requires an impressive degree of ineptitude to accidentally produce an interpretation-extending scheme that assigns both 1 and 0 to some formula.⁴

⁴ There is a subject of ‘paraconsistent logic’ where one does this on purpose, but that is not a topic for beginners (the goal is to construct logics which can tolerate contradictions, for use in situations such as database management where they will inevitably arise in the data, and one doesn’t want the entire system to become useless as a result).

2.4. Advantages of MTS as a Theory of MBPRs

As noted above, one of the things we can do with an MTS account is formulate a mathematical theory of entailment. Superficially, doing this might seem more convoluted than pursuing a deductive approach, but it turns out to have significant advantages in addition to the difficulty of actively introducing contradictions discussed immediately above. Most importantly, it is often easy to prove non-entailment, by constructing an interpretation where the premises come out true, but the conclusion false. For a very simple example, assuming the usual truth-table semantics for *or* (part of MTS) we can easily demonstrate that the following entailment does not hold:

$$(14) \quad \begin{array}{l} A \text{ or } B \\ \text{---} \text{---} \text{---} \text{or-elim} \\ A \end{array}$$

What makes MTS ‘safe’ in this respect is that it is typically relatively easy to calculate whether a sentence is true or false (or has some other relevant semantic value) for a small interpretation, so that we can check whether such an interpretation makes the premises of a deduction true but not the conclusion: if we find one, the putative entailment does not hold (which is what we want for example (14)).

Proving that entailments hold on the other hand can be a bit more difficult, since it involves mathematical arguments to the effect that any interpretation that makes the premises true makes the conclusion true as well, but these are often at the relatively modest level of difficulty of sophomore abstract algebra proofs, which people capable of getting through graduate programmes in linguistics can attain.

A potential issue is that the proofs often use natural language counterparts of the symbols and formation rules used in the deductive rules being investigated, and therefore have an appearance of circularity. We might, for example, justify the ‘soundness’ (the property of always leading from true premises to a true conclusion, relative to some choice of MTS) of the *and*-intr rule as follows:

- (15) Suppose the two premises *A* and *B* are true under *I* (i.e. *I* evaluates them as 1). Then *I* will evaluate their conjunction *A and B* as true. Therefore, since this will hold for any choice of *I*, *and*-intr is sound.

I think a close investigation of what’s really going on in proofs of this nature might be rewarding, but to begin with, one can regard them as a strategy for justifying proposed rules in a format that has been subject to very heavy testing in the last century or so of development of mathematics (when contemporary techniques for proof were developed and stabilized).

More benefits can be derived if a collection of deduction rules can be found that are provably ‘complete’ as well as sound; completeness means that every model-theoretically valid entailment also has a deductive proof. Then, one can use whichever

technique, deductive or MTS, is easiest or otherwise most suitable for the task at hand. In particular, the existence of a model-theoretic proof of soundness for an inference pattern shows that it has a formal deductive proof, even if nobody has managed to find one yet; this can be useful (Wolfgang Schwartz, p.c.).

The amelioration of the arbitrariness problem by MTS on the other hand is more ‘empirical’. In principle, it seems like there should be a stupendous variety of ways in which MTS interpretations can be set up, resulting in the same kind of bewildering variety of options that afflicts the deductive approach, but in fact, people only seem inclined to worry about a manageably small number. The fact that there isn’t a clear story about why this is the case should not inhibit us from benefiting from it.

2.5. Algebraic Semantics

MTS and deduction can function individually as alternative forms of account of MBPRs, but also jointly, in what has come to be called ‘algebraic semantics’, pioneered by Godehard Link, Erhard Hinrichs and others.⁵ In algebraic semantics, one starts with an MTS framework, and adds additional constraints that all interpretations must obey, often ones that have been extensively studied in branches of algebra, such as especially the theory of lattices and ordering. This allows the investigator to have their cake and eat it too, in the sense that they can freely propose deductive laws that appear to be true, using the MTS infrastructure to remain sure that the system as a whole is not contradictory (if it is, no sentence will have a model that satisfies it, so producing a single one that does is sufficient to demonstrate consistency).

We can for example propose that the *some of* relationship between amounts of stuff (the extensions of count and mass nouns) obeys laws such as these:

- (16) a. Transitivity: If x is some of y and y is some of z , then x is some of z .
 b. Asymmetry: If x is some of y , then y is not some of x (implying that nothing is some of itself, which seems empirically correct for the English expression *some of*).

These are the laws for a ‘strict partial order’; quite a lot is known about their consequences, and those of their variants, as part of lattice theory.

In the third section of this paper, we will discuss a number of examples of principles of this nature that appear to apply to some of the Wierzbickian primes.

2.6. ‘Support’ and Extensionality

The final useful feature of MTS that I will discuss here is its relation to the notion of ‘support’ of natural language sentences by various kinds of collections of sensory

⁵ Link (1998) contains many examples; Landman (1991) provides a textbook-style presentation.

information, and the related concept of extensionality. In textbooks, the former issue is often discussed in a somewhat limited way in connection with, for example, Cresswell's (1978) picture of a door with respect to establishing the truth or falsity of sentences such as *The door is open* and *The door is closed*. For the realist, this is a simple example of the real-world significance of language; for the antirealist, less so, because the response will in general be to a limited amount of information served up by the senses and other cognitive abilities in what might be reasonably regarded as a 'language of thought'. One might therefore treat them as a form of entailment, but since these phenomena do not involve relationships between actual, utterable, sentences, I think it is better to treat them separately, calling this second notion 'support'.

Support as characterized here, but not under this name, has been the subject of substantial psychological investigations, such as Johnson-Laird and Byrne (1991, 1993), and, for a recent example, Pietroski *et al.* (2009), which investigates hypotheses about the semantics of *most* in terms of how sensory information is accessed. Viewed as a theory of support, conventional MTS can be interpreted as the claim that many kinds of propositions, such as those involving *and* and *every*, are assessed in terms of combinations of multiple pieces of information that can be expressed in terms of atomic formulas. So, for example, in order to assess *every child is wearing a hat*, we scan the scene before us, and, whenever we notice an x such that $Child(x)$ is true, we try to find a y such that both $Hat(y)$ and $Wearing(x, y)$ are true. If all of these searches are successful, we assess the sentence as true, otherwise not (and run the risk of proclaiming a falsehood if we miss a short, hatless child lurking behind a taller one).

One interesting characteristic of support is that all of the standard deduction rules can be motivated by experiences involving support, in that they seem to lack finite counterexamples.⁶ For example, if somebody says in good faith *every crab in the bucket is green*, and you reach in and pull out one that appears to be red, it always appears to be the case that there is some 'tricky' feature to the situation. Perhaps there were many crabs in bucket and they didn't shift them around enough to find the red one hidden under the others; perhaps there is something strange about the conditions (such as UV light and gene engineered fluorescence), or some other explanation. But something more plausible than the amazing discovery that the rule of universal instantiation is wrong always seems to turn up. We can conclude from this that experience with finite model checking is plausibly the basis of many or even all entailment intuitions, which perhaps seem 'necessary' because of the extremely large amount of empirical testing that they have undergone.

Closely related to the kinds of experiences covered by the term 'support' are the phenomena of truth-functionality and extensionality. Viewed intuitively, the sentences *John loves Mary* and *Mary does not love John* are two unfortunately

⁶ This is my interpretation of a p.c. from Barbara Partee of a remark made to her by Johann van Benthem.

connected universes of feeling. But to assess the truth of their conjunction *John loves Mary and Mary does not love John*, we can ignore all of the complexities of the situations supporting the two sentences, and need to consider only the resulting truth-values, computing that of the conjunction by means of the truth-table. Similarly, with *every crab in the bucket is green*, the truth-value of the result does not depend at all on the details of the speaker or hearer's conception of crabs, buckets or the colour green, but only on the set theoretical relation between the collection of crabs in the bucket and the subset of those crabs that are green (for truth, identity is required). In the next section we will see that truth-functionality and extensionality are limited to a very small number of primes. But they are nevertheless remarkable phenomena, worthy of close investigation. Support is a notion which is not, as far as I can see, a subject that is suitable for systematic empirical investigation by linguists (or philosophers), although useful studies of specific topics can be done with the aid of psychologists. But it can be regarded as providing a 'context of inspiration' whereby reflections on the use of language in everyday situations can be used to suggest mathematical ideas which, through the techniques of model theory, can be used to construct rigorous formal accounts of the MPBRs.

2.7. Conclusion

This concludes our sketch of formal semantics, with focus on the issue of why antirealists should take MTS seriously even if they don't accept the arguments that are usually produced to support it. In the next section, we introduce NSM, as preparation for discussing how concepts of formal semantics can be applied to it, in spite of the programmatic rejection of the proclaimed realist goals of MTS that is a salient feature of the current NSM literature (e.g. Goddard 2011: 46).

3. Natural Semantic Metalanguage

The Natural Semantic Metalanguage (NSM) programme was initiated by Anna Wierzbicka in the early 1970s (Wierzbicka 1972), and has been under development by her and many colleagues ever since. The collection of NSM publications has become very large; a fundamental one is Wierzbicka (1996), while a recent major one is Goddard and Wierzbicka (2014). A introductory textbook is Goddard (2011), and the NSM homepage presently at <https://www.griffith.edu.au/humanities-languages/school-languages-linguistics/research/natural-semantic-metalanguage-homepage> provides updated information.

We start with a quick sketch, and a small sample study to give a sense of how it works (and why it is methodologically not as different from formal semantics as one might think at first), and then proceed to consider a variety of more specific topics where formal semantics has something to say about the behaviour of NSM primes.

3.1. Quick Sketch of NSM

The basic idea of the programme is in conception very different from formal semantics. It is to state meanings in a universal, intuitively intelligible manner, using a small number of words, and, more recently, constructions, found in some form (there are issues here) in all languages. This might seem completely impossible, and there certainly are problems, but it seems to me that in the mid-1980s, especially with the appearance of Wierzbicka (1985), the programme crossed a threshold in terms of being able to provide much more convincing ‘explications’, the standard term for NSM accounts of meanings, than had previously been attained, especially for words with rich meanings, such as *cat*, *dog* and *mouse*, steadily extending to more semantic classes (mostly of nouns and verbs).

The basic units out of which the explications are constructed are called ‘(semantic) primes’ (originally, ‘primitives’), each of which is supposed to have an ‘exponent’ (sometimes several, used under different circumstances) in every language. So the prime ‘because’ has exponents *because* and *because of* in English. Exponents can be multiword expressions, and primes can share exponents (BECAUSE and AFTER often do in Australian languages, for example). Although the primes are supposed to be universal, their exponents are obviously not, and of course the superficial forms of the grammatical constructions will also differ (so that the universal syntax for NSM is basically a format for ‘tectogrammatical’ structure in the sense that this term is sometimes used in the categorial grammar community, e.g. Dowty (2007: 58–59)). So NSM can be seen as an abstract universal language, essentially a kind of Language of Thought (Fodor 1987), which is furthermore manifested in a concrete subset of every natural language, yielding what can be called ‘the NSM’ of that language.

Although these basic ideas have remained constant, contemporary versions of NSM have significant differences from the earliest ones:

- (17) a. The set of primes is now much larger: around 65 instead of the original 13.⁷
- b. There is now attention to the task of disciplining the syntax as well as the word-choice in explications, so that extending the approach beyond its original and still primary target of lexical semantic to compositional semantics is no longer impossible in principle (although very underdeveloped in practice). Furthermore, the account of lexical semantics is constrained by this practice, since proposed explications have been abandoned because they use a syntactic construction without any evident rendering in some language.
- c. There is now a concept of ‘semantic molecule’, which is an intermediate-

⁷ At an average growth rate of one every year or so, there would appear to at least a century’s worth of headroom before the size of the inventory gets uncomfortably close to 200, which seems to me to be a plausible approximate maximum size for the inventory. And of course this is a proposed worst case; the prime set will not necessarily continue to grow at this rate.

level concept that is first explicated and then used in further explications. It is currently unclear whether they can be eliminated (replaced by their definitions) in individual explications or not.⁸ Molecules arguably first made their first appearance informally in Wierzbicka (1985) as a practical technique to make explications of complex concepts such as *mouse* and *lion* intelligible, and were later elevated to the status of an explicit component of the theory. Some examples of what appear to be universal molecules are *hand*, *woman* and *child*. For discussion, see Goddard (2012).

- d. In addition to word-explications, NSM is used to formulate ‘cultural scripts’, often involving the good and bad primes, which supplement many aspects of meaning that can’t be managed in a convincing way with definitions. Such scripts allow us to reconcile the apparent universality of words meaning at least something similar to English *good* and *bad* with the simultaneous existence of enormous differences in the extensions of these terms.

Molecules and semantic scripts are conceptually very important innovations, because their inclusion means that NSM no longer aspires to be a programme of merely replacing words with eliminative paraphrases stated with the primes, a goal whose feasibility has been widely, and, I think, correctly, questioned.

Instead, one may think of it as a programme for producing what are in effect teaching materials, so that if you learn the molecules, you will be able to read scripts and explications, thereby learning concepts and behavioural norms implicit in other languages and cultures with greater hope of not mindlessly importing all of the assumptions built into your own. This seems far more feasible than constructing eliminative paraphrases, and is furthermore an enterprise where ‘mere’ progress as opposed to perfection seems likely to be useful, even in real world practice.

3.2. NSM Sampler

To give a bit of the flavour of NSM, I will consider the explications of the two English causative verbs *have* and *make* provided by Wierzbicka (2006: 171–178) and Goddard (2011: 313–314), with some comparisons to the rather intellectualized verb *cause* that has often dominated discussion of semantic primitives outside of the NSM tradition. The causative verbs all involve the *because* prime, which is at least roughly equivalent to the verb *cause* when an event is its subject:

- (18) a. Mary left the room because John was sneezing.
b. John’s sneezing caused Mary to leave the room.

⁸ The programming language Prolog provides an example of predicates that can be defined but not eliminated from the places where they are used.

However, since *because* is much more characteristic of normal as opposed to intellectualized speech, appears (for example, in forms such as ‘*coz*’ in the language of two year olds (Goddard 2001)), and is, easier to use to construct explications, it is used as the prime. The verb *cause* also diverges considerably in meaning from *because* when its subject is animate. Two of the anonymous reviewers furthermore pointed out that it tends to be used when the caused event is viewed as undesirable, which is a very clear reason not to take it as the prime.

Consider first the contrast between the following:

- (19) a. Mary had John open the safe
b. Mary made John open the safe

Both seem to entail the following:

- (20) Mary did something
Because of this, after this, John opened the safe.⁹

But they differ considerably in the nature of what Mary is said to have done: example (19a) entails that there is a ‘chain of command’ from Mary to John; John might be a bank employee with Mary his superior, but she would probably not be a member of a criminal gang compelling John to open the safe so they could rob it. The reverse is the case for example (19b).

In the case of example (19a), Mary says something to somebody, but not necessarily to John, but she seems to want John to do the job, and might merely have relayed the order through an intermediary. There also seems to be an expectation that John thought he ought to comply; the first ingredient gives us a ‘chain’, the second that it is a chain of ‘command’. Goddard (2011: 314) expresses this as the following explication:

- (21) *Someone X had someone Y do something*
Someone X wanted someone Y to do something
X knew that if someone said to Y something like ‘X wants you to do this’, Y would do this because of this
Because of this, X said something like this to someone
Because of this, after this, Y did as X wanted

This explication illustrates in its second component ‘X knew that if ...’ the theoretically problematic flexibility of NSM syntax in actual use; I would suggest reformulating that portion of the explication in a more uniform format like this:

⁹ ‘Because of this, after this’ is a recurrent formula in NSM explications of verbs describing actions.

- (22) X knew this:
 if someone said to Y something like this: X wants you to do this
 Y would do this because of this

One can justify this flexibility on the basis that being immediately intelligible is a very important goal for NSM, which can at least temporarily outweigh rigorous adherence to a specific formal theory, although in the end both intelligibility and adherence to a specific theory are required. More compact formulations such as that of (21) can therefore be accepted for a while on the basis of being easier to take in, although we ultimately need to work out how to systematically reduce them to a minimum inventory of constructional forms. But in spite of the syntactic flexibility, the words in (21) are all prime exponents.

Unregulated informalities are indeed very common in NSM syntax; the goal is to progressively remove them without a sudden sacrifice of intuitive intelligibility. Another issue that is not addressed here, or very systematically in NSM to date, is exactly how the participant roles in the explication get bound to the syntactic satellites of the verb.

Wierzbicka's (2006: 176) explication is the following, somewhat less structured than Goddard's:

- (23) *Person X had person Y do Z:*
- a. X wanted something to happen (to W)
 - b. because of this, X wanted Y to do Z (to W)
 - c. because of this, X said something to someone
 - d. because of this, Y did Z
 - e. X could think about it like this:
 when I say something like this (about something like this) Y can't say 'I don't want to do this'

One difference is that the component a. 'X wanted something to happen to Y' is dropped by Goddard, who in effect elevates Wierzbicka's b. to a full clause in the explication, which captures the entailments of both a. and b. Another is that Wierzbicka's explication doesn't capture the idea that Y complies because of the chain of command, which Goddard's does. Note how Goddard's formulation in terms of know captures the mutuality of the social knowledge involved in the 'chain of command'.

Moving on to *make*, a first observation is that there are a number of uses with different grammatical environments (Goddard 2011: 315–317). Here we consider *make* with an animate subject and VP able to denote a deliberate action, which usually entails deliberate action by the subject on the object to produce the result described by the VP. Goddard's explication (2011: 316) is:

- (24) *Person X made person Y do Z:*
 X wanted Y to do something because of this, X did something

because of this, Y thought something like this: 'I can't not do this'
 Y did this because of this, not because of anything else

This explication starts out the same as the one for *have* except for omitting 'PERSON', which I take to be an insignificant variation, but completely lacks any correspondent to the second component of the *have* explication, which is significant, since the meaning of *make* seems to (correctly) entail nothing in the way of social preconditions or expectations about Y, only sufficient force or other persuasive power to get Y to comply. The second clause of (24) is like the third clause of (21), again, significantly, since *make* doesn't seem to imply any specific form of action on the part of the subject/causer, while *have* implies communication ('saying'). The last two clauses capture Y's sense of reluctance, and doing it only under compulsion, missing from the meaning of *have*. For example, if the coordinator of the park care group I am a somewhat indolent member of 'has' me occupy a publicity table at the local shopping centre, I don't do it because I don't think there's any alternative, and for no other reason, but rather because I think she wants me to (and may have other reasons, such as wanting to support the cause, even if my performance is sometimes deficient).

Wierzbicka's explication of *make* (Wierzbicka 2006: 181) is very similar to Goddard's, so I won't discuss it here, but instead say a bit about *cause*. With an animate subject and VP designating a deliberate action, *cause* seems to share something like the first two clauses of *make*, but then diverges by not attributing any reasons, attitudes or thought processes to the person who performs the desired action. I also think that in fact the first clause is a bit more general than for the other two verbs. For example, if Mary causes John to trip the secret alarm, she may only want the alarm to go off, and does something that causes somebody to trip it, without intending that that person be John. There is finally the negative attitude component, which I attribute to the speaker:

- (25) *X caused Z [=Y to do something]*
 X wanted something (Z) to happen because of this, Y did something because
 of this, Z happened
 < I think that Z was bad >

The angle brackets indicate that this final component might be optional (it is certainly absent in many academic registers, such as philosophical discourse).

The requirement of deliberateness is shown by the fact that example (26a) requires that Mary has sneezed on purpose (perhaps in an exaggerated manner) in order to drive John from the room, not just, say, from having a cold.¹⁰ With an action nominal subject, the deliberateness entailment disappears, and Mary can just be sneezing normally due to having a cold:

¹⁰ As pointed out by an anonymous reviewer.

- (26) a. Mary caused John to leave the room [by sneezing]
 b. Mary's sneezing caused John to leave the room

Cause also belongs to a more formal register than the other two verbs, more appropriate to court cases and philosophy seminars than to everyday conversation. Accounting for such sociolinguistic register effects is something that NSM could and I think should aspire to do, but relatively little attention has been directed towards it. Considerable further information about the meaning of *cause* can be found in the discussion of why *kill* does not mean *cause to die* by Fodor (1970) and Wierzbicka (1975).

4. NSM Meets Formal Semantics

In this section we work through some areas of potential contact between NSM and formal semantics, starting with a general comparison, and then considering a number of specific areas where themes from formal semantics emerge in NSM.

4.1. Theoretical Comparison

Although NSM certainly is very different from formal semantics, there is one important commonality: they both have a strong dependence on the same kind of data, intuitions of entailment (including anomaly of observed examples explained by conflicting entailments), although especially in NSM, these are supplemented by corpus data.¹¹ One difference is that in formal semantics, the entailments primarily hinge on syntactic structure, while in NSM, they tend to hinge on word choice, such as the differences caused by *make* vs. *have*. This clearly facilitates the use of corpus data, via the expectation that collocations involving contradictory or discordant entailments will be relatively rare.

But there is a more fundamental difference in how the entailments are explained by the theories. Since the formulas/abstract linguistic structures used in formal semantics are not natural language, they can only derive any kind of meaning they might have via formal apparatus associated with them¹² (unless one assumes that, with use, the notations become naturally intelligible in the way that normal languages are, which is I think sometimes assumed implicitly, but not systematically justified). This could be

¹¹ NSM publications in particular make considerably use of cooccurrence data from texts, but based on my observations of seminars and workshops, these appear to play more of a supporting role than to be the primary techniques of discovery. That this is in fact the case is indicated by the fact that neither of the words 'corpus' or 'cooccurrence' appears in the indices of Wierzbicka (2006) or Goddard (2011), although observations of natural use abound. Formal semantics can of course use corpus data in the same way that NSM does, as an additional check on hypotheses, to compensate for some of the problems with intuitions about possible use rather than observations about actual use.

¹² NSM researchers tend to believe that they have no meaning in any sense whatsoever, but that is not an issue I want to contest here.

taken as at least one aspect of the sense of David Lewis' (1972: 169) famous assertion that 'semantics with no treatment of truth is not semantics'. Namely, it is not enough just to propose some abstract notations for displaying meanings. And even for antirealists, the apparently rather central role of truth in the concept of entailment should suggest that a theory of semantics without some kind of story about truth is at best incomplete, and probably rather lame (true is in fact one of the NSM primes).

With NSM explications, the situation is different: since these are in natural language lightly supplemented with a bit of whitespace formatting similar to what is familiar and widely used in placards, menus, billboards and other kinds of public writing display, they are meaningful as they are, without requiring any additional formal work (and reference indexes, if they turn out to be required, seem to occur naturally in manual sign languages). If somebody claims that it is in general impossible to assess the entailments of an explication without a formalization, they are thereby claiming that most of the data of natural language semantics, including formal semantics, do not exist, which is absurd.

It is plausible to argue (Maïa Ponsonnet, p.c.) that this difference is not absolute, but a matter of degree: NSM is somewhat conventionalized, in ways that make it not always entirely straightforward for beginners, and people do seem to find it easy to pick up the intuitive meaning of various kinds of logic notations. But nevertheless, I think the difference in degree is rather large: full logical forms are usually far more bewildering to beginners than explications are, and the simple clauses out of which explications are constructed have a high degree of immediate intuitive intelligibility compared with logical formulas. Furthermore, Gladkova *et al.* (2015) finds no significant correlation between NSM knowledge and evaluation of explications. NSM explications thus constitute an intriguing case of an explanatory theory which doesn't need a mathematical formulation in order to have empirical consequences. Which does not, of course, imply that attempts to formalize it wouldn't lead to interesting results; and in particular, a serious but failed attempt to do formal semantics on NSM would suggest that formal semantics is impossible, and perhaps lead to some insight into why. This is I think quite a sufficient reason to investigate the prospects.

Below, I will consider two aspects of the situation. In the first place, I will review some logical properties of the primes, with attention to truth-functionality, extensionality and 'intensionality' in the sense of applicability of possible worlds semantics, as opposed to 'hyperintensionality' in the sense of Pollard (2008), as manifested with some of the mental state primes, such as THINK, SAY and KNOW, which I suggest should be seen as the expected case.

4.2. Truth-Functionality

Sheer truth-functionality is a very limited phenomenon in NSM. There appear to be only three or four words and constructions in the prime set that exhibit it, these being TRUE, NOT and possibly IF ... , along with the grammatical construction whereby multiple

sentences are combined into an explication or sub-explication, which I call ‘box composition’, which effectively performs the role of logical conjunction.

I will begin with the comparatively new prime TRUE, which seems to first appear in 2002 (Wierzbicka 2002a, 2002b; Goddard & Wierzbicka 2002), and still has received relatively little discussion. It appears to be used as a predicate to NP subjects, or as an attributive to (SOME)THING. For example, in Wierzbicka (2002a: 411), the explication of the Russian sentence in example (27) starts out as example (28):

- (27) Ivan skazal nepravdu
Ivan told/said not.truth
Ivan told an ‘untruth’
- (28) Ivan said something
He knew that this
<something> was not true

The angle brackets indicate that the ‘something’ word here is sometimes included in such explications, sometimes not, for reasons that could use further reflection. Omitting it does not make the explication ungrammatical or unintelligible. (And keep in mind that the full explication is complex and specific to Russian, and doesn’t necessarily correspond closely to the explications of English words such as *lie*, *false(hood)*, *untrue* or *untruth*—the last being relatively rare, but nevertheless existing.)

Truth-functionality reveals itself in discourses such as:

- (29) Susan understated her tax liability. This is true.

The antecedent of *this* is the first clause, and the second clause is true if and only if the first clause is. Truth-functionality likewise obtains if the second clause is changed to *this is not true*, but the discourse becomes contradictory, whereas, in the original version, the first clause seems to function as an initial admission that is expected to lead up to some kind of plea for indulgence, such as *but nobody ever told her that her grandmother had set up a Swiss bank account in her name*.

In English we can also have:

- (30) a. it is (not) true that Susan understated her tax liability
b. That Susan understated her tax liability is (not) true

Following Wierzbicka (2002b), I suggest that this is not a valence of the true prime, but is rather to be explicated along the lines of:

- (31) Some people maybe think like this:
Susan understated her tax liability
This is (not) true

This attempts to capture the idea that the construction is not primarily used to add a new assertion to the common ground, but to comment on something that is already in it, at least as a possibility.

There is on the other hand no prime ‘false’; English *false* is not a prime exponent. To see this, consider the oddness in ordinary discourse of something like:

- (32) I thought that Cindy was in her office. This was not true/?false

The negation of *true* is banal here, *false* seems strange. I suggest that this is because the meaning of *false* involves components that refer to rhetorical postures whereby the non-truthfulness of the assertion is claimed to have some kind of wider significance:

- (33) I used to think that Poverty of the Stimulus arguments motivated a Morphological Blocking Principle. Now I believe that this is false.

A tentative starter explication of *false* might run along the following lines:

- (34) *X is false*
 Some people think like this:
 X
 This is not true
 It is good if people know this

We don’t here account for the fact that the first component appears to be a presupposition rather than an assertion.

We have already seen uses of NOT together with TRUE, and it is clear that its meaning flips truth (not ‘truth-values’, in NSM *per se*, because MTS, where the concept of truth-values resides, is not intrinsically part of NSM, but rather a tool that might prove useful for investigating NSM). We can ask if NSM is bivalent in the sense that every sentence is either true or false, and the answer presumably has to be ‘no’, because of the paradoxes. Interestingly, many of the explications and cultural scripts in Wierzbicka (2002a) involve formulations of bivalence as one of their ingredients, indicating that it is not being thought of as a basic features of NSM, but something that needs to be stated explicitly when needed. For example, the full explication of example (27) is:

- (35) Ivan said something
 He knew that this something was not true
 people can say two kinds of things to each other things of one kind are true
 it is good if people want to say things of this kind to other people
 Ivan did not say something of this kind things of the other kind are not true

it is not good if people want to say things of this kind to other people
Ivan said something of this other kind

Note that although this explication implies that *true* and *not true* are important kinds of things that people can say to each other, it does not entail that there are not others, such as questions, and paradox sentences that can be neither true nor not true, such as *What I am saying now is a lie*, given plausible ideas about what kinds of entailments hold, and that *true* and *not true* are mutually exclusive.

Aside from bivalence, the other logic-based question about negation would be the status of double negation, that is, does *not not S* entail *S*. Due to the restrictions we've imposed on the valence of TRUE, we can't produce something close to a standard double negation formula without introducing extraneous elements, but can contemplate a discourse like this:

- (36) A: John is not washing the dishes.
B: That is not true!

This appears to entail that John is washing the dishes, arguing that double negation is valid for NSM, and therefore for ordinary language.¹³

To progress further with formalizing truth-functionality in NSM, it might seem most natural to proceed deductively, but, as discussed in the previous section, the model-theoretic approach seems easier. It would appear to at least be convenient (perhaps there are other ways) to have at least two 'truth-values', conventionally either {1, 0} or {T, F} (as noted above, the non-existence of a false prime does not imply that we can't use a 'false' truth-value in an MTS for NSM, since MTS can be regarded as a tool for investigating MBPRs in NSM, rather than an intrinsic part of NSM). So then, clauses in an explication will at least sometimes get evaluated as 1 if they are true, 0 if they are not true (and the logic might be 'three-valued' with an 'unspecified' value for some combinations, paraconsistent with some sentences having both T and F values, and many other possibilities). We then propagate these values to instances of this that are anaphoric to those clauses, propagate them unchanged to predications with true, and flip them if NOT is present. To fill out this sketch as a proper formal theory, it would be necessary to provide an account of anaphora, which will not be attempted here.

The next truth-functional construction is simply the operation whereby one clause is added to one or more others to form a multiclausal explication. Logically, this is simply conjunction (in spite of frequent NSM assertions to the effect that *and* is not the exponent of a prime), and has the truth-table of logical conjunction. In particular, if an explication is true of a situation, every component in it must be true, and if every component in an explication is true, whatever it is that it explicates

¹³ And therefore that intuitionistic logic, useful as it is, does not provide an accurate representation of ordinary language meaning.

must be true, if the explication is correct. In order to proceed deductively here, we would need to figure out some way to deal with the anaphoric references, which seems like it might be challenging, but evaluating anaphors to the reference of their antecedents, which we have already discussed, and treating box composition as having the truth-conditional semantics of conjunction would be a plausible way to start.

Our remaining at least partly truth-conditional construction is IF The truth conditional aspect is revealed by the obvious falsity condition for example (37), where the *if*-clause (protasis) is true and the main clause (apodosis) false:

(37) If Avery's door is open, he is in his office.

The problem case is the classical one whereby for the 'material conditional', the whole thing is true if the protasis is false. This debate has raged for literally millenia ('Even the crows on the roof are cawing about which conditionals are true', Callimachus); a relatively recent survey is Bennett (2003).¹⁴

We obviously have no chance whatever of settling this issue here, but we can see the problem by contemplating this conversation:

(38) A: If Avery's door is open, he is in his office.
 B: ?*Yes indeed, it's closed
 B': ?Yes indeed, it's open and he's there

The intuition here is that to 'support' a conditional we need evidence for a general rule that lacks exceptions, for which answer B is completely irrelevant, and B' is insufficient, constituting only a single case. On the other hand, the claim that the conditional is fully truth-functional is supported by uses like this:

(39) If that guy knows anything about biology, then I'm the Pope [spoken by somebody who's clearly not the Pope, perhaps, Jerry Coyne]

The idea here is that the obvious falseness of the apodosis means that in order for the whole sentence to be true, the protasis must be false, so the net effect is to assert that that guy does not know anything about biology. Perhaps this is an idiosyncratic feature of English discourse, but, the cross-linguistic semantics and use of conditionals is not a very well-developed area.

A final case that is not quite truth-functional, but relevant anyway, is the *when* construction. Consider a sentence such as:

¹⁴ With respect to which I note that it would be very interesting to see how many of the complex and problematic scenarios discussed by Bennett work out in the same way for the standardly recognized conditional constructions of other languages, especially non-European ones without significant cultural heritage from Ancient Greece.

- (40) When Avery's door is open, he's in a good mood.

A simple MTS account of sentences like this can be given by having the semantic values of sentences be functions from time-periods to truth-values. We evaluate this sentence by looking at all the (relevant) time periods, and checking for each the material conditional 'Avery's door is open Avery is in a good mood'. If there are no time periods for which the conditional is false, and, perhaps at least some where its protasis is true, we judge that the sentence is true (my intuition is that the sentence is not so clearly true if the door is never open). So although this construction goes beyond mere truth-functionality, requiring quantification over times in the formulation of the model theory, it does seem to involve a major truth-functional component.

4.3. Extensionality

Extensionality is in essence truth-functionality adapted to collections, or 'sets',¹⁵ of people and things. The clearly extensional words are the 'quantifiers' ONE, TWO, SOME and all. There are two other quantifiers, MUCH/MANY and LITTLE/FEW, which might also well be extensional, but involve complex contextual effects that I do not want to look at here. The basis of extensionality is that many Natural Language nouns, verbs and adjectives, in typical situations, pick out sets or sets of n -tuples of entities that they are 'true of' (that is, you can create lists of things or arrays of things that they are true of). So in a given back-yard scene, there might be entities (things that you can point to saying *this*) that are birds and others that are dogs, and also some that are barking, and also certain 'ordered pairs' (two entities, with designated 'first' and 'second'), such that the first is chasing the second, the first is watching the second, etc.

Extensionality refers to the fact that in sentences such as those below, involving the some and all primes, the truth of the whole depends only on the patterns of overlap between the sets of entities for which the properties and relations hold:

- (41) a. Some (of the) dogs are barking
 b. All of the dogs are barking
 c. Some (of the) dogs are chasing a bird

¹⁵ There is a terminological subtlety here: the term 'set' in mathematics implies commitment to some propositions such as that two sets are identical if they have the same members, and there is a set with no members (and therefore only one such), and also that sets can be members of other sets. In typical situations of use, these commitments don't create problems. But it is not clear that we really need them for natural language semantics, and we perhaps could avoid them by talking of 'collections' instead, as often happens in introductions to category theory. How to explain this and other aspects of mathematical discourse in NSM is an interesting topic, but not one that I will attempt to pursue here.

The richness of whatever the concepts of birdhood, dogginess, barking or chasing are is completely irrelevant; only the overlaps matter.

So, for example, in example (41a), the *dog* set must have a non-empty overlap with the *bark* set, while in example (41b) it must be a subset of the bark set (the presence or absence of *of the* determines aspects of how the dog set is circumscribed, which I won't consider here). The transitive example is too complex to work through here, but is dealt with in introductions to formal semantics or symbolic logic, and it is clear that what matters is what *n*-tuples are in the set associated with *chase*, and what members in the sets associated with the nouns, rather than anything else about their meanings. Extensionality is therefore inherently surprising for the same reason that truth-functionality is. A number of related properties of 'quantifiers' were presented in Barwise and Cooper (1981), and have been mainstays of work in formal semantics ever since.

4.4. Intensionality and Hyperintensionality

'Intensionality' can be regarded as a weakened form of extensionality, including truth-functionality.¹⁶ To get a sense of what is involved, consider a sentence such as:

- (42) It is possible that Bond might be in Zurich.

The truth of this is clearly independent of whether Bond actually is in Zurich. Rather, what seems to be involved is the plausibility of Bond's being in Zurich given other knowledge. This is a delicate matter. With the arrival of the Boeing 707, it became possible for somebody to be almost anywhere in the world with real roads within 48 hours of their being spotted in any urban centre, but, still, a sincere utterance of example (42) seems to require a bit more than this level of plausibility; roughly, I think, enough of a probability to be worth deploying resources to cope with Bond's presence in Zurich in case he's really there.

The conventional way to deal with this kind of intuition in formal semantics has been 'possible worlds', developing some proposals of Carnap (1947). In contemporary implementations, an interpretation now contains a set of objects called 'possible worlds', alongside the set of entities, and a 'proposition' is a function from possible worlds to truth-values. Sentences are furthermore evaluated 'at' a possible world, so that *Bond is in Zurich* can be evaluated to 1 in some possible worlds, 0 in others (or the same in all), and is viewed as true if it evaluates to 1 in the chosen world we're evaluating at (thought of as representing the 'real world'). 'Possible', whose explication involves the MAYBE prime, plus perhaps additional ingredients,¹⁷ is then a

¹⁶ There is terminological variation here; I am following Pollard (2008).

¹⁷ See Wierzbicka (2006: 277) for a proposed explication, which does not however seem to me to capture the 'plausibility' constraint that I suggested at the beginning of the section. This explication appears in an extended discussion of many related works, such as likely, presumably, etc.

‘modal operator’ which produces a proposition that evaluates to 1 in the world of evaluation if it evaluates to 1 in one of the ‘possible’ worlds, perhaps subject to further restrictions such as ‘accessibility’ (for which, see a textbook on modal logic). We can have multiple sets of possible worlds useful for different sorts of things, so that if example (42) were to appear in a real conversation concerning an actual operative called Bond, the ‘possible worlds’ would be those deemed worth allocating some resources to deal with this in case they turn out to be the real one.

All of this might seem rather pointless until we get to the next grade, hyperintensionality (Pollard 2008). A hyperintensional operator is one which can produce results of different truth-value when applied to logically equivalent operands. For example, consider the case of Fred, who is blind, not well educated, not very smart, and perhaps faces some other ill-understood cognitive and perceptual challenges. Fred thinks that he was abducted by aliens last night, and that they have told him that they will destroy the earth if he is not holding an equilateral triangle at 7:30AM this morning (too early to recruit a friend in time to help). He is for various reasons unable to perceive equilaterality with his senses, but thinks he remembers that equiangular triangles, which he can detect, are also all equilateral, but is not quite sure.

Now suppose we have independent knowledge of the alien’s activities, about which Fred is in fact correct. We can now say either:

- (43) a. Fred must be holding an equilateral triangle at 7:30AM this morning or the earth will be destroyed.
 b. Fred must be holding an equiangular triangle at 7:30AM this morning or the earth will be destroyed.

(On the assumption that the local geometry of space is close enough to Euclidean that the relevant theorems hold in it, up to the aliens’ ability or desire to make the measurements precise.) This is because we know that the two types of triangle are the same, so that (a) will be true in all the relevant ‘possible worlds’ that (b) is and vice-versa.

By contrast, the following are not equivalent for us:

- (44) a. Fred knows that he is holding an equilateral triangle.
 b. Fred knows that he is holding an equiangular triangle.

It seems perfectly possible that (a) could be true and (b) false, the same if we replace *knows* with *thinks*, and even more so if we replace it with *says*. So, these are hyperintensional rather than just intensional. There are numerous theories about how to manage hyperintensionality in formal semantics; we won’t try to go into this here, but merely note that it seems to be a relevant distinction, which effects a classification of the primes.

In particular, it would appear that *MAYBE* and *CAN* are clearly intensional, and also, even if *IF...* is not truth-functional, it is at least intensional rather than

hyperintensional, even in the counterfactual conditional construction of (b) that we don't discuss here:

- (45) a. If Fred is holding an equiangular/equilateral triangle at 7:30AM, the earth will not be destroyed.
 b. If Fred had not been holding an equiangular/equilateral triangle at 7:30AM, the earth would have been destroyed.

Likewise, BECAUSE, HAPPEN and DO, illustrated for the first:

- (46) Because Fred was holding an equilateral/equiangular triangle at 7:30 AM, the earth was not destroyed.

On the other hand, THINK, SAY, KNOW and WANT are hyperintensional. Since FEEL has only a nominal, not a propositional valence, its status is currently unclear to me.

So we now have the intensional/hyperintensional classification applying to 10 primes, along with the three clearly truth-functional ones, which is beginning to constitute a noticeable fraction of the prime list. One might still be disdainful of these concepts on the basis that they don't apply to many primes, but this would be wrong, because the real surprise is that they apply to any at all. Meanings are, after all, very complex, and our default expectation should be that the semantic behaviour of combinations of primes should depend on all of the semantic properties of the primes that are combined. But this turns out not to be the case: in a significant number of instances, much of the meaning of a prime can be discarded for ascertaining the truth of the combinations that it appears in, which is a strange and noteworthy phenomenon.

The final topic I will look at is 'laws', essentially the application of algebraic semantics to NSM.

4.5. Laws

A major role in mathematical reasoning is played by algebraic principles which relationships may or may not obey. The most important ones that people have found it useful to study are:

- (47) a. Transitivity: For all x, y, z , if xRy and yRz , then xRz (e.g. *richer than*)
 b. Reflexivity: For all x , xRx (e.g. *identical to*)
 c. Irreflexivity: For all x , not xRx (e.g. *richer than*)
 d. Symmetry: For all x, y , if xRy , then yRx (e.g. *similar to*)
 e. Asymmetry: For all x, y , if xRy , then not yRx (e.g. *richer than*)
 f. Antisymmetry: For all x, y , if xRy and yRx , then $x = y$ (e.g. *least as much as* for amounts; this law has relatively few convincing NL examples, although it is very useful in mathematics).¹⁸

Relations obeying combinations of these (and other) principles have various interesting properties (see Partee *et al.* (1993) for an introduction oriented towards the needs of linguists). So-called ‘equivalence relations’, for example, a feature of introductory abstract algebra, obey Transitivity, Symmetry and Reflexivity, and have the property of splitting up a set exhaustively into non-overlapping subsets called ‘equivalence classes’ (*has the same birthday as*, over people, is an example).

We can now see that many of the relational primes obey various of these laws, although there are also unclear cases, such as whether NEAR is reflexive. Here are some examples:

- (48) a. PART: Asymmetric and Irreflexive, possibly Transitive (not entirely clear)¹⁹
 b. THE SAME: Reflexive, Symmetric, Transitive (arguably the prototypical equivalence relation; all the others are ‘the same in some respect’)
 c. OTHER: Irreflexive, Symmetric, (not Transitive)
 d. MORE: Transitive, Asymmetric (and therefore Irreflexive)
 e. BEFORE: Transitive, Asymmetric (and therefore Irreflexive)
 f. AFTER: Transitive, Asymmetric, (and therefore Irreflexive)
 g. ABOVE: Transitive, Asymmetric, (and therefore Irreflexive)
 h. BELOW: Transitive, Asymmetric, (and therefore Irreflexive)
 i. INSIDE: Transitive, Asymmetric, (and therefore Irreflexive, assuming that something can’t be inside itself; if we challenge this, then it because Antisymmetric and Reflexive)
 j. NEAR: Symmetric, certainly not Transitive, probably Irreflexive (this tends to produce disagreement in the classroom)
 k. FAR: Symmetric and Irreflexive; neither Transitive nor Reflexive
 l. LIKE: Symmetric; not Transitive but arguably Reflexive²⁰

Presumably the real reason for the obedience of the primes to the various laws is their connection to the human sensory apparatus and thereby to the world; it is also worth noting that the ultimate basis of most of the mathematics is thinking about sensory experience without a very close dependence on the overt form of ordinary language. So further investigations of Primes and Laws would very likely be rewarding. And MTS with the ‘algebraic semantics add on’ would seem to be a safe and promising way to investigate this.

¹⁸ Note the subtlety that least as tall as is not antisymmetric; two people can be at least as tall as each other but still be different people. Their heights however will be the same.

¹⁹ Amongst the issues here is the nature of the difference between *part of X* (e.g. a car) and *a part of X*.

²⁰ The issue of the Reflexivity of like arises with sentences such as ‘John isn’t like John, he IS John’; there is an intuition that this is odd, but that can be plausibly explained by the Gricean Maxim of Quantity (based on a p.c. from Andy Egan).

5. Conclusion and Further Issues

We have seen various ways in which NSM and formal semantics including model theoretic semantics are at least somewhat compatible in their aims and methods. Although there are many problems remaining to be solved in order to fit them together as truly complementary methods for investigating the same subject-matter, the discussion here shows that some progress can be made. Of the problems needing attention, I suggest that amongst the most important ones are coreference between multiple participants, which NSM proposes to handle without using variables, without having made concrete suggestions about the most problematic cases, and the semantics of colour and other unimodal sensory terms, where ostension appears to be required to get any substantive account of meaning, and attempts by NSM researchers to avoid this do not appear to be convincing.

The most important practical ingredient supplied by formal semantics that is lacking in NSM is mathematical formulation; however NSM has its own special feature that we have not given much explicit discussion of, which we can describe as ‘immediate intuitive intelligibility’. While entire NSM explications can be difficult to take in, it is important to note that the individual components are not. Somebody who claims to have difficulty in understanding the meaning of *this is good* or *somebody did something bad to someone* can, I think, be dismissed as hopeless just as surely as someone who rejects a simple syllogism. It should be interesting to find out whether mathematical development and immediate intuitive intelligibility can be combined into a single theory of meaning in natural language.

References

- Barwise J & R Cooper 1981 ‘Generalized quantifiers and natural language’ *Linguistics and Philosophy* 4: 159–219.
- Barwise J & R Cooper 1983 *Situations and Attitudes* Cambridge MA: MIT Press.
- Barwise J & J Perry 1985 ‘Shifting situations and shaken attitudes’ *Linguistics and Philosophy* 8: 105–161.
- Bennett J 2003 *A Philosophical Guide to Conditionals* Oxford: Oxford University Press.
- Carnap R 1947 *Meaning and Necessity* Chicago IL: University of Chicago Press.
- Chierchia G & S McConnell-Ginet 2000 *Meaning and Grammar: an introduction to semantics* 2nd edition Cambridge MA: MIT Press first edition 1990.
- Cresswell MJ 1978 ‘Semantic competence’ in F Guenther & M Guenther-Reutter (eds) *Meaning and Translation: philosophical and linguistic approaches* London: Duckworth. pp. 9–27.
- Dowty D 2007 ‘Compositionality as an empirical problem’ in C Barker & P Jacobson (eds) *Direct Compositionality* Oxford: Oxford University Press. pp. 23–101.
- Fodor J 1970 ‘Three reasons for not deriving “kill” from “cause to die”’ *Linguistic Inquiry* 4: 429–438.
- Fodor J 1987 *Psychosemantics* Cambridge MA: MIT Press.
- Gladkova A, U Vanhatalo & C Goddard 2015 ‘The semantics of interjections: an experimental study with natural semantic metalanguage’ *Applied Psycholinguistics*. doi:10.1017/S0142716415000260
- Goddard C 2001 ‘Conceptual primes in early language development’ In M Pütz & S Niemeier (eds) *Applied Cognitive Linguistics I: theory and language acquisition* Berlin: Mouton de Gruyter. pp. 193–227.

- Goddard C 2011 *Semantic Analysis: a practical introduction* 2nd edition Oxford: Oxford University Press.
- Goddard C 2012 'Semantic primes, semantic molecules, semantic templates: Key concepts in the NSM approach to lexical typology' *Linguistics* 50(3): 711–743.
- Goddard C & A Wierzbicka 2002 *Meaning and Universal Grammar: theory and empirical findings* Vols I and II Philadelphia PA: John Benjamins.
- Goddard C & A Wierzbicka 2014 *Words and Meanings: lexical semantics across domains, languages, and cultures* Oxford: Oxford University Press.
- Jackendoff RS 1983 *Semantics and Cognition* Cambridge MA: MIT Press.
- Jackendoff RS 1985 'Information is in the mind of the beholder' *Linguistics and Philosophy* 8: 22–33.
- Johnson-Laird PN & RM Byrne 1991 *Deduction* Hillsdale NJ: Lawrence Erlbaum Associates.
- Johnson-Laird PN & RMJ Byrne 1993 'Précis of deduction' *Behavioral and Brain Sciences* 16: 323–380.
- Katz JJ 1966 *The Philosophy of Language* New York: Harper and Row.
- Katz JJ 1972 *Semantic Theory* New York: Harper and Row.
- Katz JJ & J Fodor 1964 'The structure of a semantic theory' In *The Structure of Language: readings in the philosophy of language* Englewood Cliffs NJ: Prentice-Hall. pp. 479–518 first published in 1963, *Language* 39:170–210.
- Keenan EL & LM Faltz 1985 *Boolean Semantics for Natural Language* Dordrecht: Reidel.
- Kneale W & M Kneale 1962 *The Development of Logic* Oxford: Oxford University Press.
- Lakoff G 1970 'Linguistics and natural logic' *Synthese* 22: 151–171. Available at http://deepblue.lib.umich.edu/bitstream/handle/2027.42/43819/11229_2004_Article_BF00413602.pdf?sequence=1
- Landman F 1991 *Structures for Semantics* Kluwer Dordrecht: Academic Publishers.
- Larson R & G Segal 1995 *Knowledge of Meaning* Cambridge MA: MIT Press.
- LePore E 1983 'What model theoretic semantics cannot do?' *Synthese* 54: 167–187.
- Lewis D 1972 'General semantics' In D Davidson & G Harman (eds) *Semantics of Natural Language* Dordrecht: D. Reidel. pp. 169–218.
- Link G 1998 *Algebraic Semantics in Language and Philosophy* Stanford CA: CSLI Publications.
- Montague R 1970a 'English as a formal language' In B Visentini (ed.) *Linguaggi nella Società e nella Tecnica* Milan Rome: Edizioni di Comunità. pp. 189–224 reprinted in Thomason (1974). pp. 188–221.
- Montague R 1970b 'Universal grammar' *Theoria* 36: 363–398 reprinted in Montague (1974).
- Partee BH 2004 'Reflections of a formal semanticist' In *Compositionality in Formal Semantics: selected papers of Barbara Partee* Oxford: Blackwell Publishers longer version available as Partee (2005). pp. 1–25.
- Partee BH 2005 *Reflections of a formal semanticist* Available at http://people.umass.edu/partee/docs/BHP_Essay_Feb05.pdf accessed 27 January 2015 (longer version of Partee (2004)).
- Partee BH, A ter Meulen & RE Wall 1993 *Mathematical Methods in Linguistics* Dordrecht: Kluwer Corrected second printing.
- Pelletier FJ 1999 'A brief history of natural deduction' *History and Philosophy of Logic* 20: 1–31.
- Pietroski P, J Lidz, T Hunter & J Halberda 2009 'The meaning of "most": Semantics, numerosity, and psychology' *Mind and Language* 24: 554–585.
- Pollard C 2008 'Hyperintensions' *Journal of Logic and Computation* 18: 257–282 Prepublication draft available at: <http://www.ling.ohio-state.edu/~pollard/cvg/hyper.pdf> accessed 27 January 2015.
- Thomason R 1974 *Formal Philosophy: selected papers by Richard Montague* New Haven CN: Yale University Press.
- Wierzbicka A 1972 *Semantic Primitives* Frankfurt: Athenäum Verlag.
- Wierzbicka A 1975 'Why "kill" does not mean "cause to die": The semantics of action sentences' *Foundations of Language* 13: 491–528 reprinted in Wierzbicka (1980).

- Wierzbicka A 1980 *Lingua Mentalis: the semantics of natural language* London: Academic Press.
- Wierzbicka A 1985 *Lexicography and Conceptual Analysis* Ann Arbor, MI: Karoma.
- Wierzbicka A 1996 *Semantics: primes and universals* Oxford: Oxford University Press.
- Wierzbicka A 2002a 'Russian cultural scripts: The theory of cultural scripts and its applications' *Ethos* 30: 401–432.
- Wierzbicka A 2002b 'Semantic primes and universal grammar in Polish' In A Wierzbicka & C Goddard (eds) *Meaning and Universal Grammar: theory and empirical findings* Vol I Philadelphia PA: John Benjamins. pp. 65–144.
- Wierzbicka A 2006 *English: meaning and culture* Oxford: Oxford University Press.