1 *Trivium pursuits*

> But the truth is, they be not the highest instances that give the securest information, as may be well expressed in the tale so common of the philosopher that while he gazed upwards into the stars fell into the water; for if he had looked down he might have seen the stars in the water, but looking aloft he could not see the water in the stars. So it cometh often to pass that mean and small things discover great, better than great can discover the small. Bacon, *The Advancement of Learning*, Book II, 1.v. (1605)

1.1 As above, so below

Bacon's philosopher might be forgiven for looking too much upwards and not enough down. We look "up" not just to the stars and the sky, but to those we admire and to our highest ideals. We look "down," as often as not, on things we despise, things beneath us, which are low, mean, and base. Familiarity breeds contempt, and it is easy to forget that what lies beneath may also run deep.

Figuratively speaking, up is where it's at. *Up* is above, on top of, superior to, beyond; it is higher than, taller than, farther than, and more. It can be a location or a direction. It is defined within a larger frame, the vertical scale, which it shares with *down* – normally, the physical dimension parallel to an upright person standing erect on an even surface. The basic experience of bodily uprightness motivates the common metaphorical associations of being "up" with wakefulness, alertness, strength, reason, and virtue, and being "down" with sleep, weakness, folly, and vice. This massive alignment of evaluative metaphors along a vertical scale is not just some whim of imaginative fancy, nor is it unique to English. Indeed, it is a normal way for conceptual contents to be imaginatively structured across semantic domains – a reflection in grammar of the workings of the mind.

The basic opposition between 'up' and 'down,' and the many metaphorical oppositions it engenders, are themselves symptoms of a much more general tendency for human concepts to be structured in terms of contraries. All languages, it seems, have metaphors in which abstract notions like 'truth' and

1

2   *The Grammar of Polarity*

'goodness' are fleshed out in terms of more basic bodily experiences, and one of the most basic experiences featured in such metaphors is the sense of opposition one may feel between contrary concepts like 'up' and 'down,' 'light' and 'dark,' or 'hot' and 'cold.' Contrariety itself is a quintessentially abstract concept, but it is immanent in our most down-to-earth experiences. The human mind thrives on the logic of contraries, and this is everywhere reflected in the structure of language, from the most basic phonemic oppositions and antonymic lexical pairings to the elementary rules for predicate affirmation and denial.

Keeping with Bacon's advice, this book looks mainly down at little things in order to glimpse therein the image of something great. The little things of concern here are matters of grammar – ordinary constructions of everyday talk and their attendant bits of form and meaning. The greater things to be discovered are the elements and principles of thought itself: the commonsense imaginative abilities which allow us, the speaking ape, to entertain concepts and to share them with one another.

1.2   A quirk of grammar or a trick of thought?

This book is concerned with a single, intricate, and easily overlooked grammatical phenomenon going by the awkward name of *polarity sensitivity*. Many, and perhaps all, human languages include a class of constructions which are somehow sensitive to the expression of polarity – forms whose acceptability in a sentence can depend on whether that sentence is grammatically negative or affirmative. Such *polarity items* arise in many semantic domains and come in many morphosyntactic flavors; but, since polarity itself is a binary relation, all polarity items divide into two basic classes: positive polarity items (PPIs), which are unacceptable in the scope of negation, and negative polarity items (NPIs), which are unacceptable in simple affirmative contexts.

Both NPIs and PPIs can be found side by side in semantic domains they share with semantically similar but grammatically insensitive (or *neutral*) constructions. The data in (1–4), for example, reveal four sets of sensitivity triplets – items with similar semantics but different sensitivities – taken from four basic semantic domains: (1) agentive effort, (2) epistemic possibility, (3) propositional conjunction, and (4) event frequency. For each domain, the examples in (i) illustrate neutral items, those in (ii) illustrate PPIs, and those in (iii) illustrate NPIs. The unacceptable sentences in (ii–iiib) give some impression of what happens when a polarity item occurs in the wrong sort of context.

(1)   EFFORT: (i) *make an effort to* V, (ii) *take a stab at* V-*ing*, and (iii) *even bother to* V.
    i)   a.   He *made an effort* to solve the puzzle.
        b.   He didn't *make an effort* to solve the puzzle.
    ii)   a.   He *took a stab* at solving the puzzle.
        b.   *He didn't *take a stab* at solving the puzzle.
    iii)   a.   *He *even bothered* to solve the puzzle.
        b.   He didn't *even bother* to solve the puzzle.

(2)   POSSIBILITY: (i) *be likely to* V, (ii) *could well* V, and (iii) *can possibly* V.
    i)   a.   She *is likely to* win the race.
        b.   She *is* not *likely to* win the race.
    ii)   a.   She *could well* win the race.
        b.   *She *could*n't *well* win the race.
    iii)   a.   *She *can possibly* win the race.
        b.   She *can*'t *possibly* win the race.

(3)   CONJUNCTION: (i) *and*, (ii) *as well as*, and (iii) *let alone*.
    i)   a.   Chris has read the Aeniad *and* the Georgics.
        b.   Chris hasn't read the Aeniad *and* the Georgics.
    ii)   a.   Sally has read the Aeniad *as well as* the Georgics.
        b.   *Sally hasn't read the Aeniad *as well as* the Georgics.
    iii)   a.   *Glynda has read the Aeniad, *let alone* the Georgics.
        b.   Glynda hasn't read the Aeniad, *let alone* the Georgics.

(4)   FREQUENCY: (i) *to V X a lot*, (ii) *be always* V-*ing* X, (iii) *to V X much*.
    i)   a.   Ann listens to the Grateful Dead *a lot*.
        b.   Ann doesn't listen to the Grateful Dead *a lot*.
    ii)   a.   Hugh *is always* listen*ing* to the Grateful Dead.
        b.   *Hugh *is*n't *always* listen*ing* to the Grateful Dead.
    iii)   a.   *Jeff listens to the Grateful Dead *much*.
        b.   Jeff doesn't listen to the Grateful Dead *much*.

The proper way to account for this little phenomenon has been a subject of long-standing and at times rather intense controversy in theoretical linguistics. These are not the sorts of facts one is likely to notice about a language, but they are remarkable nonetheless. One would expect that anything one could affirm, one could also deny, and that anything one could deny, one could also affirm. But polarity items are subject to special constraints, the violation of which results in unexpectedly unacceptable sentences. These constraints are more complicated than the examples here suggest since NPIs can be licensed, and PPIs blocked, in a variety of contexts beside clausal negation – among others, in questions, and in conditional (*if*) and comparative (*than*) clauses (see below §2.3.2). Still, the fundamentally striking observation here is that a simple switch in polarity can make an otherwise unobjectionable sentence

not just unacceptable, but apparently ungrammatical. The problem with these
sentences is not just one of semantic anomaly (since it is clear what they should
mean) nor of any obvious pragmatic infelicity (for it is easy to see how they
might be used). Rather, something about these sentences seems to make them
intrinsically incoherent. The question is, what is the nature of this incoherence?
What, precisely, is wrong with these sentences? How should this wrongness
be represented in a theory of grammar? And crucially, what is it about the way
speakers understand such sentences that makes them feel so wrong?

To answer these questions, one must confront fundamental questions about
the nature of grammar and meaning. Almost from the start of generative lin-
guistics, polarity items have been a battleground in debates about the nature
of grammatical representation (Lees 1960; Bolinger 1960; Klima 1964; Baker
1970), and as theories have evolved, polarity items have remained a flashpoint.
Polarity sensitivity neatly straddles the realms of syntax, semantics and prag-
matics, so that a theory of polarity necessarily raises questions not just about
the interfaces between these components, but ultimately about the architecture
of grammar itself and the grammar's relation to extra-linguistic aspects of cog-
nition (Fauconnier 1975a, 1976; Ladusaw 1979, 1983; Linebarger 1980, 1987;
Israel 1996, 1998a, 2004; Chierchia 2004; Giannakidou 2006). For the most
part, these debates have turned on the question of what sorts of entities are
needed in a theory of grammatical representations in order to account for the
constraints on polarity items.

The distributions of polarity items have thus served as evidence that the
grammaticality of a sentence may depend on its entailments (Baker 1970)
or on its implicatures (Linebarger 1980, 1987, 1991), and as such they have
played a central role in debates about the nature of logical form as a level in
grammatical representations. Most famously, perhaps, Ladusaw (1979, 1983)
has argued that the grammar of polarity items depends on a fully interpreted
level of logical form where negative polarity items are constrained to appear in
the immediate scope of a downward entailing (DE) operator. According to this
proposal, the model-theoretic representation of a sentence's literal truth condi-
tions is itself a part of grammar – a level where constraints on well-formedness
are defined – and not merely the product of more general cognitive abilities
operating on the output of a generative grammar.

However one chooses to formulate the constraints on polarity items, one
must also confront the problem of how language users manage to learn these
constraints. Polarity items epitomize a classic quandary of language acqui-
sition: the absence of negative evidence (Braine 1971; Bowerman 1988;
Pinker 1989). Somehow speakers learn the grammar of polarity items without

hearing the ways these forms cannot be used. But what speakers have to learn about polarity items is precisely the ways they are not used. The obvious way one could learn such a thing – the way linguists in fact learn it – is to find an instance of a polarity item in a context where it cannot be used and to observe – whether by introspection or controlled elicitation – the oddness of its usage. But of course ordinary speakers can never make such an observation since the oddness, or "ungrammaticality," of such uses normally prevents their occurring at all.

In fact, one of the few places such uses do occur (though even here they are rare) is in the spontaneous speech of very young children. The examples below, from the CHILDES database (MacWhinney 1995), illustrate the sort of uncertainty typical of young children's early uses of polarity items. In (5–6), Abe is just under 33 months old (2;8.22), when he uses the idiomatic NPI *in my life* in a conversation with his father about an orange fish (Kuczaj 1976):

(5)   *FAT: I bet if you used one of those orange fish # you could catch something what do you think?
           *ABE: what orange fish?
           *ABE: what orange fish?
           *ABE: I never heard of that my life.
           *FAT: you never heard of that in your life?
           *ABE: I wan(t) (t)a go catch a corn fish.   (File 032 – lines 47–53)

In this first use, the NPI (or something close to it, since Abe actually omits the preposition *in*) is licensed by the negative *never*. Most likely Abe has learned the NPI here [*in* pro*'s life*] as part of a larger idiom – something like *never heard of X in my life*. But whatever the details, Abe's usage here is clearly flexible and creative, as moments later he produces the same item in a simple affirmative sentence, without *never* or any other negative licensor.

(6)   *FAT: what kind do you want to catch?
           *ABE: a [/] a [/] a [/] a stair fish.
           *FAT: a stair fish?
           *ABE: uhhuh I heard of that in my life.
           *FAT: you heard of that is [*sic*] your life?
           *ABE: uhhuh I can't fish like that.   (File 032 – lines 118–23)

Apparently, Abe at this age was not yet aware of the constraints which limit expressions like *in my life* to negative contexts, or if he was, he did not yet realize that this particular expression is subject to such constraints.

A similar pattern of confusion appears in Nina's corpus (Suppes 1974), where, on one occasion at 36 months (3;0.16), the child seemed

6   *The Grammar of Polarity*

to vacillate between *any more* and *some more* in several repetitions of the same clause.

(7)   *NIN: have to close (th)em # (be)cause it's not raining any more.
         *NIN: when it's raining some more.
         *NIN: it's not raining some more now.
         *NIN: it's not raining any more # so we have to close this one.   (File 44 – lines 640, 652, 653, and 687)

If these sorts of anecdotal observations are at all representative (and they are certainly not uncommon), it appears that whatever children might know about the theoretical constraints on polarity items, it is not enough to keep them from using such items in some very unconstrained ways.

Even if one assumes that speakers come equipped with some innate knowledge of the constraints which govern polarity items, speakers still must learn the particular constructions in their language that are sensitive, which sensitivities they have, and just how strongly sensitive they are. This is a formidable problem since languages vary widely both in the polarity items they include and in the details of their distributions. Moreover, as the data in (1–4) show, near synonyms can and do vary sharply in their sensitivities. Somehow, it seems, speakers must master these subtleties on a case-by-case basis. It thus seems reasonable to follow van der Wouden's suggestion (1997: 80), that while "the mechanisms underlying the behaviour of polarity items are part of grammar; the specific behaviour of individual polarity items is part of the lexicon." Still, the question is, just how do these grammatical mechanisms find their way into the individual polarity items?

This book seeks answers to this and other questions about the grammar of sensitivity by viewing polarity items, and sensitive items in general, in terms of the semantic and pragmatic contents they encode in observable discourse (whether "real" or in some way experimentally contrived). I assume, in other words, that polarity items are polarity sensitive because of the meanings they encode, so that speakers effectively learn "the grammar" of these constructions (i.e. their particular sensitivities) the same way they learn the meaning and use of any other linguistic construction.

This does not mean that "the grammar" here is not in some sense "innate" or "universal." There are universal constraints on what a human mind may imagine, and on what sorts of imaginings can be encoded by a linguistic construction. But such constraints might take a variety of forms, and it is far from clear which, if any, of our innately human predispositions consists precisely in a constraint on linguistic representations. I will argue here that the distributions

of polarity items, at least, are not determined by constraints on linguistic rep-
resentations per se, but rather reflect the operation of general cognitive abilities
in ordinary communicative interactions.

My goal is to explain not just why polarity items have the peculiar distribu-
tions they do or how speakers manage to learn these distributions, but also why
it is that polarity items should exist in the first place. I argue that polarity sen-
sitivity in general arises as a grammatical consequence of the ways language
users regularly exploit a basic conceptual ability for rhetorical purposes. The
conceptual ability here is the ability to reason in terms of scales – the ability,
that is, to construe an entity within a particular sort of semantic frame, a scalar
model, and to make inferences based on this construal.

1.3   The hypothesis: sensitivity as lexical pragmatics

The basic theory – what I call the *Scalar Model of Polarity* – is simple. The
claim is that polarity contexts are defined by their effects on scalar inferences
and that polarity items encode semantic properties which make them sensitive
to such inferences. Polarity items are thus a special class of what Fillmore,
Kay, and O'Connor (1988) and Kay (1990) term "scalar operators" – forms
which must be interpreted with respect to an appropriately structured scalar
model.

In particular, I claim, sensitivity arises from the interaction of two sorts of
scalar semantic properties – quantitative (q-) value and informative (i-) value –
each of which functions independently of polarity sensitivity, but which
together constitute the necessary and sufficient conditions for a construction to
be polarity sensitive. A form's q-value depends on its relative position (either
high or low) in a scalar ordering. A form's i-value reflects the informative
strength (either emphatic or attenuating) of the proposition to which the form
contributes its meaning. Both features are grounded in the logic of scalar rea-
soning and the rhetoric of interpersonal communication. Their combination
within a single form effectively limits that form to contexts which allow the
scalar inferences needed to make both values felicitous.

The theory makes clear predictions about where polarity items might be
found in a language and what forms they can take. Most generally, the theory
predicts the existence of four broad classes of polarity items: NPIs divide into
emphatic forms with low q-value and attenuating forms with high q-value;
PPIs divide into attenuating forms with low q-value and emphatic forms with
high q-value. All four sorts are well attested in English and other languages,

and the theory predicts that all sorts of polarity items from all sorts of domains fit this broad taxonomy.

The idea that sensitivity might be related to scalar semantics is not new: it has been advanced in one way or another by an impressive set of theorists (e.g. Schmerling 1971; Horn 1972, 1989, 2005; Fauconnier 1975a, 1976; Fillmore, Kay & O'Connor 1988; Kadmon & Landman 1993; Lee & Horn 1994; Krifka 1995; Haspelmath 1997; Lahiri 1998; van Rooy 2003; Zepter 2003), and disputed in one way or another by an equally impressive set (Linebarger 1980; Progovac 1992, 1994; Rullmann 1996; Giannakidou 1998, 1999; Chierchia 2004; Szabolsci 2004). The present work, however, makes the unusual claim (though see Verhagen 2005 for a similar view) that polarity items are not just scalar in their propositional semantics, but also in their pragmatics.

Polarity items are, I contend, argumentative operators[1] which conventionally index an argumentative attitude – an attitude, that is, toward the expressed content of an utterance; or, in Gricean terms, toward what is (baldly and explicitly) said. For my purposes here, it will suffice to distinguish just two major types of argumentative attitude, *emphasis* and *attenuation*, each of which may attach to either a positive or a negative proposition. Constructions which express an emphatic attitude – for example, the English [*really* Adj] and [(*not*) *at all* Adj] constructions in (8) and (9) – present an expressed proposition (what is said) as somehow stronger and more significant than an alternative proposition which might have been said. Conversely, constructions expressing an attenuating attitude – like [*sort of* Adj] and [(*not*) *such a* Adj] in (10) and (11) – hedge what is said, and present a proposition as weaker and less exciting than it might have been.

(8)	a.	That's true.	p
	b.	That's really true.	p (> n)
(9)	a.	That's not true.	~p
	b.	That's not true at all.	~p (> n)
(10)	a.	That's a good idea.	q
	b.	That's sort of a good idea.	q (< n)
(11)	a.	That's not a good idea.	~q
	b.	That's not such a good idea.	~q (< n)

The constructions here illustrate the four basic sorts of argumentative meanings. These are very general sorts of meaning, and as such can be (and typically are) coded by a great many constructions within a single language. The notations on the right reflect the status of these sentences as neutral, emphatic,

or attenuating: "p" and "q" here stand for expressed propositions, "(n)" for a salient alternative proposition (the parenthesis indicates its status as implicit or backgrounded), and the "more than" (">") and "less than" ("<") signs show the strength of an expressed proposition relative to n, and thus its status as either emphatic or attenuating.

There are a variety of ways one might understand "strength" as a property of propositions – as, for example, its likelihood of being true (Karttunen and Peters 1979), its noteworthiness (Herburger 2000), its relevance (van Rooy 2003), or its force as an argument for some conclusion (Ducrot 1973, 1980; Anscombre and Ducrot 1983). I follow Kay (1990, 1997) in defining the strength of a proposition directly in terms of its entailments: a proposition p is stronger than a proposition n if and only if p unilaterally entails n. I take it that while emphasis and attenuation are fundamentally rhetorical aspects of meaning, they are in fact grounded in this simple propositional logic. Marking an expressed proposition as either emphatic or attenuating is basically just a way of calling attention to its logical status with respect to background assumptions. But the act of calling attention itself is always rhetorically loaded. An argumentative operator thus does not add to the logical content of what is said but expresses an attitude about that content and so situates it in a larger context.

The key idea in this book is that such argumentative content is an irreducible part of the meanings of certain linguistic constructions, and that the encoding of such content has systematic grammatical consequences. This idea seems to go against the grain of much contemporary theorizing. The problem is that emphasis and attenuation are fundamentally pragmatic aspects of meaning, and so the claim that sensitivity depends on such features means that polarity licensing must be, at least in part, pragmatic in nature. The "grammaticality" of a polarity item in a linguistic context is thus a function not just of the sentence in which it occurs, but also of the utterance in which it is used. But if this is true, and if sensitivity really is a grammatical phenomenon (which, I maintain, it is), then grammar itself cannot be limited to the generation of well-formed sentences but must also regulate their uses in discourse.

Whether or not this claim really makes sense depends in part on how one imagines linguistic knowledge (or grammar) is mentally encoded and how it relates to communicative competence in general. Basically, it makes sense if, as I contend, pragmatics is a part of grammar, and linguistic constructions regularly encode pragmatic constraints as an irreducible part of their conventional meanings. It makes less sense if one assumes that grammars are strictly a matter of linguistic representations, that pragmatics merely effects the ways such representations can be used, and that pragmatic effects are, as a rule,

dependent on objective semantic contents of linguistic constructions. Both sorts of assumptions have much to recommend them, so it may be useful to consider some of the reasons why the latter view is so widely assumed, and why it thus seems so odd to so many that grammar might be in some measure a matter of pragmatics.

1.4   Putting pragmatics in its place

It was Charles Morris who in 1938 first distinguished pragmatics as a branch of semiotics distinct from and parallel to the studies of syntax and semantics. In some ways, Morris's trichotomy of syntax, semantics, and pragmatics is reminiscent of the original trivial pursuits, the Trivium of grammar, logic, and rhetoric – the first three of the seven liberal arts. Carnap famously conceived of this trichotomy as a series of abstractions:

> If in an investigation explicit reference is made to the speaker, or, to put it in more general terms, the user of a language, then we assign it to the field of *pragmatics*… If we abstract from the user of the language and analyze only the expressions and their designata, we are in the field of *semantics*. And if, finally, we abstract from the designata also and analyze only the relations between the expressions, we are in (logical) *syntax*.   (Carnap 1942: 9)

Thus syntax – or at least "(logical) syntax" – studies relations between logical or linguistic expressions; semantics is what you get when you add meanings to those expressions; and pragmatics is what you get when you place those meanings in contexts with speakers, hearers, and communicative intentions. Framed this way, syntax – the combinatorics of signs – is at the (logical) core of the enterprise. One advantage of this is that it immunizes the study of linguistic forms from the obvious subjectivity which infects so much of language use. At the same time, by isolating syntax from both meaning and usage, it presupposes that the forms of syntax are, in a deep sense, independent from the communicative concerns of the people who use them.

This way of framing things – in particular the idea that grammar, meaning, and use belong to separate domains in the study of language – has been a cornerstone of generative linguistic theorizing. Generative grammar starts from the assumption that a language is a kind of formal object – "a set (finite or infinite) of sentences finite in length and constructed out of a finite set of elements" (Chomsky 1957: 13) – and that a grammar is "a device that generates all of the grammatical sequences of [a language] L and none of the ungrammatical sequences" (*ibid*.). As Partee, ter Meulen, and Wall put it (1990: 437),