

The Baltic International Yearbook of Cognition,
Logic and Communication

Volume 3: *A Figure of Speech*

August 2008
pages 1-47

MICHAEL GLANZBERG
University of California, Davis

METAPHOR AND LEXICAL SEMANTICS

ABSTRACT: This paper shows that several sorts of expressions cannot be interpreted metaphorically, including determiners, tenses, etc. Generally, functional categories cannot be interpreted metaphorically, while lexical categories can. This reveals a semantic property of functional categories, and it shows that metaphor can be used as a probe for investigating them. It also reveals an important linguistic constraint on metaphor. The paper argues this constraint applies to the interface between the cognitive systems for language and metaphor. However, the constraint does not completely prevent structural elements of language from being available to the metaphor system. The paper shows that linguistic structure within the lexicon, specifically, aspectual structure, is available to the metaphor system.

This paper takes as its starting point an observation about which sorts of expressions can receive metaphorical interpretations. Surprisingly, there are a number of expressions that *cannot* be interpreted metaphorically. *Quantifier expressions* (i.e. *determiners*) provide a good example. Consider a richly metaphorical sentence like:

- (1) Read o'er the volume of young Paris' face, And find delight writ there with beauty's pen; Examine every married lineament (*Romeo and Juliet* I.3).

In appreciating Shakespeare's lovely use of language, *writ* and *pen* are obviously understood metaphorically, and *married lineament* must be too. (The meanings listed in the Oxford English Dictionary for *lineament* include diagram, portion of a body, and portion of the face viewed with respect to its outline.) In spite of all this rich metaphor, *every* means simply every, in its usual literal form. Indeed, we cannot think of what a metaphorical interpretation of *every* would be.

As we will see, this is not an isolated case: while many expressions can be interpreted metaphorically, there is a broad and important group of expressions that cannot. Much of this paper will be devoted to exploring the significance of this observation. It shows us something about metaphor. In particular, it shows that there is a non-trivial *linguistic* constraint on metaphor. This is a somewhat surprising result, as one of the leading ideas in the theory of metaphor is that metaphor comprehension is an aspect of our more general cognitive abilities, and not tied to the specific structure of language.

The constraint on metaphor also shows us something about linguistic meaning. We will see that the class of expressions that fail to have metaphorical interpretations is a linguistically important one. Linguistic items are often grouped into two classes: *lexical* categories, including nouns, verbs, etc., and *functional* categories, including determiners (quantifier expressions), tenses, etc. Generally, we will see that lexical categories can have metaphorical interpretations, while functional ones cannot. This reveals something about the kinds of semantic properties these expressions can have. It also shows that we can use the availability of metaphorical interpretation as a kind of *probe*, to help distinguish these sorts of categories.

Functional categories are often described as 'structural elements' of language. They are the 'linguistic glue' that holds sentences together, and so, their expressions are described as being semantically 'thin'. Our metaphor probe will give some substance to this (often very rough-and-ready) idea. But it raises the question of whether all such structural elements in language—anything we can describe as 'linguistic glue'—are invisible when it comes to metaphorical interpretation. We will see that this is not so. In particular, we will see that linguistic structure that can be found *within* lexical items may be available to metaphorical interpretation. This paper will show specifically that so-called *aspec-*

tual structure, which language uses to encode certain ways of depicting events in a linguistically structured way, is available to the cognitive system responsible for metaphor interpretation.

The picture that will emerge from this discussion is that the cognitive system involved in metaphor is highly selective in how it accesses linguistic structure. It can uniformly ignore certain structural elements, including functional categories; while at the same time it can access other structural elements, like aspectual structure.

This paper will proceed as follows. Section 1 will show that determiners cannot be interpreted metaphorically. Section 2 will generalize this observation, to show that while the lexical categories can be interpreted metaphorically, the functional categories cannot. This section will also explore the ways this allows us to use metaphor as a probe for investigating the linguistic notion of functional category. Section 3 will turn to the question of how to understand the constraint on metaphor our observation about functional categories indicates. It will argue that it is best understood as a constraint on the interface between the cognitive systems involved in language and metaphor. Section 4 will explore the extent of this constraint. It will show that some aspects of linguistic structure are available to the metaphor system. In particular, features of linguistic structure within the lexicon are. Finally, section 5 will offer some conclusions and speculations.

1. SOME MISSING METAPHORS

I shall begin by showing that certain categories of expressions lack metaphorical interpretations. This section will focus on determiners. In section 2 we will see that all expressions falling under functional categories lack metaphorical interpretations.

Along the way, I shall try to remain as neutral as possible about the notion of metaphor itself. We will have clear enough examples to know it when we see it, and will not need a fleshed out theory of metaphor to identify and work with the constraints this paper will propose. However, we will need to be a little more clear on just what we mean by a part of a sentence having or lacking a metaphorical interpretation. Before presenting the main evidence about determiners, I shall pause to clarify this point. (I shall return to more general issues about metaphor

briefly in section 3.)

1.1. Metaphors Derived from Parts of a Sentence

Some metaphors may be localized to particular parts of a sentence. We can see this in the stock example of a metaphor in philosophical discussion—the much overused:

(2) Juliet is the sun.

In interpreting this sentence metaphorically, we typically understand *Juliet* as meaning Juliet. It contributes its literal meaning to the metaphor. But we do not understand *sun* as simply contributing its literal meaning. Rather, it triggers a metaphorical interpretation involving being radiant, a source of warmth, etc. In understanding the whole sentence (or an utterance of it) metaphorically, we thus localize the metaphorical content to *sun*, but take *Juliet* literally. In such a case, as I shall say, *sun* receives a metaphorical interpretation, while *Juliet* does not.¹

A few provisos are in order about how I shall use this notion of metaphorical interpretation, mostly to point out that I mean something rather weak. First, I am not supposing that there is any such thing as a special species of metaphorical meanings that attach directly to linguistic items (or denying it either, for that matter). All I need to suppose is that expressions in a sentence play an important role in our coming to comprehend a metaphor, and that in some cases, that role sets up a mapping from an expression to its literal meaning, and in some cases to something very different from its literal meaning.²

Second, I make no claim about what the primary unit of metaphorical interpretation is, and particularly, no claim that it is subsentential. It is compatible with what I have noted that it is whole utterances, in contexts, that are the primary bearers of metaphorical interpretation (and this appears to be the standard assumption). I have only noted that when we do interpret an utterance, we can often localize metaphorical interpretation to parts of the sentence uttered, and isolate some parts that contribute only their literal interpretations.

To reinforce this, it is worth noting that whether or not a part of a sentence is interpreted literally or metaphorically is dependent on the context. Consider an example from Stern (2000):

- (3) A revolution is not a matter of inviting people to dinner. (Attributed to Mao Tse-Tung.)

If indeed Mao said this, it is likely that it was to be understood with *revolution* taking its literal meaning. (For Mao, *revolution* really meant revolution!) The locus of metaphorical content is in the phrase *inviting people to dinner*. If I say the same thing to my fiancée while planning a dinner party, this latter phrase would be interpreted literally, and presumably *revolution* would require a metaphorical interpretation.

1.2. Metaphorical Interpretations of Major Lexical Categories

Now that we have a notion of the metaphorical interpretation of a part of a sentence, we can ask what sorts of expressions have metaphorical interpretations.

A quick glance shows that we can readily find metaphorical interpretations across all the major lexical categories. These are the categories of nouns, verbs, and adjectives (adverbs being assimilated to the category of adjectives). Here are some examples, including more poetic metaphors (usually drawn from Shakespeare) and more conventional ones.

- Nouns:

- (4) a. Juliet is the *sun* (*Romeo and Juliet* II.2).
b. Girls are *sugar and spice*.

- Verbs:

- (5) a. His mother's death *hit* him hard.
b. He *hardened* his heart.
c. He that depends Upon your favours *swims* with fins of lead (*Coriolanus* I.1).

- Adjectives:

- (6) a. The painting is *blue*.
b. That idea is *transparent*.

- c. My salad days, When I was *green* in judgment: *cold* in blood (*Anthony and Cleopatra* I.5).

Many more examples may be found in any extensive list of metaphors, such as that in Lakoff and Johnson (1980). We easily find a variety of metaphors localized to the major lexical categories.

1.3. Determiners Lack Metaphorical Interpretation

When we look at the major lexical categories, we find metaphorical interpretations easily. But there are some categories where we find none. I have already mentioned the example of determiners. These are expressions like *every*, *some*, *most*, *few*, etc., that express quantification, and are interpreted as generalized quantifiers in standard semantic theory.³

Determiners do not receive metaphorical interpretations. Consider a few examples:

- (7) a. Thou, best of dearest and mine only care, Art left the prey of every vulgar thief (*Sonnet* 48).
b. It is the star to every wandering bark (*Sonnet* 116).
c. Read o'er the volume of young Paris' face, And find delight writ there with beauty's pen; Examine every married lineament (*Romeo and Juliet* I.3, repeated from 1).
d. When the blood burns, how prodigal the soul Lends the tongue vows. These blazes, daughter, Giving more light than heat, extinct in both (*Hamlet* I.3).

In each of these cases, we have a rich metaphor. But in each case, the determiner is *not* interpreted metaphorically. Each occurrence of *every* in (7) means every, literally. The same for *more*. These examples are typical of what we see when we go hunting for metaphorical interpretations for these expressions. We can find lots of metaphors in which determiners appear, but none where they themselves get metaphorical interpretations.

I have been careful to put this claim in terms of the grammatical category of *determiner*, rather than more loosely in terms of quantifier expressions. This turns out to be important, both to get the claim right, and for a generalization we will explore in a moment.

The claim I am advancing is that there are no metaphorical interpretations of determiners. This needs to be distinguished from a few others for which it might be mistaken. First, the claim is that there are no *metaphorical* interpretations of determiners. There are other figurative interpretations which are possible. We can find irony or over- or under-statement (or if you like the traditional names, hyperbole and meiosis) with determiners. For instance:

- (8) a. Like, MOST things he says are true.
b. I saw every famous philosopher at the conference.

This does require us to have a clear enough sense of what counts as a metaphor to distinguish it from other sorts of figures. In these sorts of cases, though, that is not too demanding. We see in (8) cases where the quantifying determiner (*most* or *every*) functions with its basic meaning intact. In (8a) the utterance is then interpreted as conveying the opposite of what it normally expresses. In (8b), the over-stated usage is generated by picking a stronger quantifier than is really appropriate. The utterance conveys something like that the speaker saw some famous philosophers, with an emphasis generated by the exaggeration. In none of these cases do we provide anything like a metaphorical interpretation of the determiner itself.

Second, we do find metaphorical interpretations of quantifier-like expressions when they appear in non-determiner positions. For instance, we have a metaphorical interpretation of *every* in:

- (9) He is (the) *every* man.

In this case the expression *every* is coerced into a non-determiner position. Though this is acceptable here, such coercion does not appear to be fully productive. For instance, substituting *some* or *most* for *every* gets us markedly bad sentences.⁴ There are also a few expressions, like *most*, that occur both as determiners and non-determiners. For instance, we have:

- (10) a. Most heavy day (*Anthony and Cleopatra* IV.14).
b. Make the most of it.

(I shall not try to say what the relation between the determiner and non-determiner occurrences is.) These may be apt for metaphorical

interpretation of *most*, though I am not sure. It seems more natural in these cases to simply read *most* as expressing being great in degree, which is one of its literal meanings. But either way, it is not occurring as a determiner.

It might be tempting to attribute the lack of metaphorical interpretations to something about our conceptual resources. Perhaps, one might speculate, we simply lack the conceptual resources to construct metaphorical content around whatever meanings determiners have.

This is not the case. Conceptually, quantifying determiners measure sizes. Take *most*, for instance. *Most As are B* is true if and only if $|A \cap B| > |A \setminus B|$. A related size measure is given by *More As than Bs are C*, which holds if and only if $|A \cap C| > |B \cap C|$. Yet we can readily find metaphors attached to expressions comparing sizes in just these ways. Here are a few.

- (11) a. Sure he that made us with such large discourse (*Hamlet* IV.4).
b. I do invest you jointly in my power, Preeminence, and all the large effects (*Lear* I.1).
c. It lends a lustre and more great opinion, A larger dare to our great enterprise, Than if the earl were here (*Henry IV, Part I* IV.1).
d. It will wear the surplice of humility over the black gown of a big heart (*All's Well that Ends Well* I.3).
e. Thy words, I grant, are bigger, for I wear not My dagger in my mouth (*Cymbeline* IV.2).

Metaphorical interpretations of terms like *bigger* and *larger* are commonplace.

We thus have the conceptual ability to produce metaphors with just the kinds of size comparisons and measurements that quantifiers make. Metaphors are not to be found when, and only when, the size comparison is expressed by a determiner. Conceptually, there is little difference between *larger* and the determiners *more* or *most*, but metaphor is available with the former and not the latter.

So far, I have noted that we can make sense of localizing metaphorical interpretation to a part of a sentence. When we do, we find no metaphorical interpretations of *determiners*, i.e. quantifier words. We

do find metaphorical interpretations of expressions with the same conceptual content, but not the determiners themselves.

A claim that there are no metaphorical interpretations is a hard one to make, and I should admit that I know of no fully reliable test for the lack of metaphor. But, in this case, the evidence is strong nonetheless. First of all, informally, no one has suggested any cases where we have any hint of a metaphorical interpretation of a determiner. We do not seem to have any idea what such an interpretation would be like. A little more formally, we may take Shakespeare's corpus as a rich and varied source of metaphors. Scanning a concordance to Shakespeare shows 528 occurrences of *every* and 919 of *more . . . than*. I have found none that appear to be metaphorical. With due caution, this supports the claim that metaphorical interpretations of determiners are absent. It is also striking that metaphorical uses of other size-measuring expressions are found all over the Shakespeare corpus. When we look for metaphorical interpretations of size comparison, we find them easily, but when we look for metaphorical interpretations specifically of determiners, we find none.⁵

2. FUNCTIONAL AND LEXICAL CATEGORIES

We have seen that determiners lack metaphorical interpretations. We have also seen that it is the grammatical category of determiner that seems to lack such interpretations, not the conceptual category of size-measuring phrase. That leads naturally to the question: what is so special about determiners?

Determiners themselves are not all that special, but they fall within a well known and important linguistic division between kinds of expressions. Here are some of the marks of the class of determiners. It is a *closed class*, i.e. it is virtually impossible to add determiners to a language (except by the glacial mechanisms of language change), as opposed to the open classes like nouns and verbs, where new expressions are added easily and often. Determiners do not figure into derivational morphology, i.e. there is nothing like *every-ing* or *every-ize*. They also have a highly restricted range of meanings: determiners express a certain subset of the available size-comparison operators. The restrictions on their meanings hold in every language. Again, no such restrictions

apply to the meanings of nouns and verbs.

These are the marks of what are called *functional categories*. The list of functional categories typically includes not only determiners, but also tenses (and modals and auxiliaries), complementizers (*that* and *which*), conjunction, and negation. Some theories include prepositions, and some have much longer lists. These are all, like determiners, closed classes. They are often described as those categories that provide the 'grammatical glue' that binds words together into sentences. This is perhaps most natural with respect to expressions like complementizers, which seem to serve simply the grammatical function of introducing subordinate clauses. It is also the case for tenses, modals, and auxiliaries, which provide the frame around which the rest of the linguistic items in a sentence are arranged. It is an increasingly common view that this sort of frame-providing role is common across all the functional categories.⁶

Another feature often proposed for functional categories is that they have either no semantic content, or have semantic content that is somehow thin. For instance, Fukui (2001, p. 392) says that they "do not have substantive content" while Baker (2003, p. 87) says they lack "rich distinctive lexical semantics." The idea is that functional categories primarily provide 'grammatical glue', while other sorts of expressions provide the substantial content of a sentence.

This might seem plausible for complementizers like *that* and *which*, which seem to play a grammatical role only, and do not contribute anything interesting to the meaning of a sentence. But philosophers and formal semanticists will find it strange to say that determiners and tenses lack any 'rich distinctive' semantics. Indeed, if anything, the semantics of determiners is one of the more well-understood areas of semantics. A number of interesting properties of the semantics of determiners are known, which support important linguistic generalizations, including some very likely semantic universals. This, it would seem, is as rich as we could ever want semantics to be. Much the same can be said of other operators, like tenses, which also display rich semantic properties.

To be fair, the quotes above are from works in *syntax*, and they are not really concerned with these aspects of semantics. Even so, we have something of a puzzle. Functional categories do have rich semantic properties, and yet, there does seem to be something distinctive about

their meanings. We see this when we think about the relation of their meanings to the grammatical roles functional categories play, and the limited range of meanings they can have. We thus would like to know what it is that makes the meanings of functional categories different, and why those meanings often seem ‘thin’ (at least to syntacticians).

On this question, metaphor turns out to be of some help. It shows us one way in which functional categories lack a ‘rich semantics’ which lexical categories have. We have already seen that the lexical categories all have metaphorical interpretations, while the determiners do not. This generalizes to all the functional categories. By running the same sorts of tests we ran with determiners, we see clearly that complementizers lack metaphorical interpretations. So do negation, conjunction, and related expressions (again, putting aside irony and over-statement). Tenses show exactly the same results as determiners. We have lots of metaphorical interpretations of expressions for time, but none for tenses appearing in their grammatical positions. We have lots of metaphorical interpretations of *the past*, but none for *-ed*.⁷

We thus find that metaphor can tell us something about the lexical/functional distinction:

Having metaphorical interpretations correlates with lexical categories, while lacking metaphorical interpretations correlates with functional categories.

This correlation gives some substance to the idea that functional categories have only ‘thin’ meanings; at least, insofar as metaphor is an instance of markedly ‘rich’ meaning, the lack of it may well count as ‘thin’. Furthermore, the correlation shows that we can use metaphor as a *probe* for lexical versus functional status. Probing for metaphorical readings of expressions amounts to probing for lexical versus functional status.⁸

Let us take stock of our results so far. We have learned something about metaphor, and something about lexical and functional categories. We have seen that all the lexical categories can be interpreted metaphorically, but the functional ones cannot. This shows us something about the nature of functional categories, and why their meanings are often described as ‘thin’. It also offers us a probe—a tool—for investigating lexical semantics and the lexical/functional distinction. We have thus used metaphor as a way to explore some important linguistic ideas.

We have also seen something important about metaphor itself. We have seen a substantial constraint on what can be interpreted metaphorically. Though we have identified this constraint, in terms of the linguistic notions of functional and lexical category, we have yet to explore what the constraint really shows us about metaphor. The rest of this paper will begin this exploration. It will address two sorts of questions: how a constraint derived from linguistic categories may affect metaphor, and what the extent of the constraint is. In particular, this paper will explore in more detail what aspects of linguistic structure are available for metaphorical interpretation.

3. THE METAPHOR AND LINGUISTIC SYSTEMS

We will now begin to explore the issues of how our constraint works to affect metaphor, and what the extent of the constraint really is. To do so, we will need to rely on some (I believe rather minimal) assumptions about both language and metaphor. In particular, we will need some assumptions about how language and metaphor are cognized, and how the cognitive systems involved interrelate.

In this section, I shall briefly review some assumptions about metaphor and about language. I shall then use them to construct a very simple model of how metaphor and language relate. This model will give us a picture of how our constraint on metaphor from functional categories operates, and what the extent of the constraint is. This model will be a tool for investigation, rather than a final conclusion. I shall argue in section 4 that it is wrong in important respects. Doing so will help us to better understand the scope of the constraint on metaphor. Thus, we will at least begin to address the questions before us.

The main idea we will need, both about language and about metaphor, is that the two are cognitively different. This is widely assumed. For instance, I take it that when Lakoff (1993, p. 203) says “In short, the locus of metaphor is not in language at all, but in the way we conceptualize one mental domain in terms of another,” he takes himself to be advancing a well-established thesis. Of course, metaphors are expressed using language, but whatever cognitive systems allow us to generate and understand metaphorical interpretations are assumed to be different from those systems specific to language. In the next two

subsections, I shall elaborate this idea, by filling in some background assumptions that support it.

3.1. *Assumptions about Metaphor*

In the preceding sections, I relied only on our ability to recognize metaphors and to localize them to parts of sentences. Beyond that, I assumed little about metaphor. In what follows, I shall continue to be as non-committal as possible about theories of metaphor. In particular, I shall not take a stand on any of the main points of contention in either the philosophy or psychology literatures. Even so, I shall make a few minimal assumptions about metaphor, and a brief review of these issues in philosophy and psychology will help to make those assumptions clear. This will help to make clear the differences between metaphor and language in cognition.

In philosophy, much of the debate has focused on whether metaphorical ‘meanings’ are the same as other sorts of meanings, i.e. if they are ‘propositional’ in nature. A diverse group of philosophers argue for propositional meanings, including Searle (1979), and Stern (2000). The view that metaphors have distinctive non-propositional meanings is championed by Black (1962), while Davidson (1978) argues that metaphors have no distinctive meanings at all. So long as we can recognize a difference between metaphorical and non-metaphorical interpretations of sentences and their parts, any of these is compatible with what I say in this paper.

Another of the main issues in philosophy is the mechanism by which metaphors are conveyed. Some, such as Searle (1979), take metaphor to be a (Gricean) pragmatic process whereby a speaker literally says one thing but conveys a distinct speaker’s meaning. Relevance theorists such as Carston (2002) and Sperber and Wilson (1998), and other contextualists like Recanati (2004), assimilate metaphor to the category of ‘loose use’, or more generally to the same sorts of pragmatic mechanisms they see as involved in determining what proposition is conveyed by a speaker. (They thereby downplay or reject any distinction between metaphorical and literal.) Another option, suggested by Walton (1993), is that metaphors are conveyed via a kind of pretense or make-believe game in which speakers can engage.⁹

Any of these provide for a significant non-linguistic component of

metaphor. For a speaker to convey and a hearer to understand a metaphor requires them to rely upon Gricean or relevance or pretense processes. We may safely assume that any of these go beyond linguistic competence. Ability to engage in make-believe, or to recognize relevance, or informativeness, or perspicuity, all go beyond our competence with language. The same goes for semantic views of metaphor, like that of Stern (2000). Stern sees metaphor as the result of a context-dependent operator, which has distinctive semantic properties. But the value of the operator in context is still set by pragmatic processes which go well beyond linguistic competence.

I thus think it is safe, and not too controversial, to assume that there is an important difference between what goes into metaphor and what goes into linguistic competence. Some psychological abilities and processes are required for the conveying and interpreting of metaphor that go beyond competence with words and phrases.

Assuming there is such a difference, we might ask what the cognitive processes involved in metaphor interpretation are like. The Gricean and relevance-theoretic views do tell us something about these processes, as they hold that they are essentially the same processes that are involved in other aspects of communication. It is natural to conclude that they place the cognition of metaphor within general intelligence, as presumably that is where our abilities to maximize information or optimize informativeness across a wide range of subject-matter are to be found. The same can be said for Stern’s semantic view, which relies on such cognitive resources as our ability to judge salience. Views like Black’s, which rely on a notion of metaphorical meaning, likewise see these meanings as arising from systems of concepts interacting in some ways. Again, we may assume that whatever cognitive processes go with generating or using such systems are not linguistic in nature. All of these views place the cognitive processes used to interpret metaphor squarely outside of linguistic competence, and in the realm of highly general cognitive abilities.

Many of these philosophical accounts of metaphor are proposed as highly idealized models, or as rational reconstructions, of our ability to understand and convey metaphors. Even so, they do at least strongly suggest some predictions about the genuine real-time processing of metaphor. For instance, Searle’s Gricean picture predicts that the inter-

pretation of metaphor goes through the literal meaning of an utterance; a prediction that has been much-debated in the psychology literature.¹⁰ Though I do not really want to dwell on the details of the psychology of metaphor, a brief glance at the leading ideas there will help substantiate the difference between linguistic competence and metaphor cognition.

The leading psychological models of metaphor comprehension all rely in some form or another on our abilities to work with concepts: to compare them, to generalize and abstract from them, etc.¹¹ For instance, one leading model (the *category-transfer model* in the terminology of Camp, 2006) starts with a given concept and abstracts from features of a prototypical instance of it, to form an ‘ad hoc category’. Such a category is a general schema of concepts associated with the prototype. This category is then transferred to a target concept, typically a literal content. So, for instance, Glucksberg and Keysar (1993) discuss the metaphor *My Job is a jail*, which they see involving the formation of an ad hoc category from a stereotypical instance of being a jail. Such a category includes being unpleasant, confining, punishing, etc. This category is then applied to the concept provided by *my job*, to produce the metaphor.

Other models rely on our ability to compare concepts, and find shared features between them (*feature-matching models*, in the terminology of Camp, 2006). These models often rely upon our ability to judge *salience*, to restrict metaphor interpretation to salient shared features. Ortony (1979) suggests that metaphor is marked by an imbalance in salience, where the shared features are highly salient with respect to one concept, but not the other. Finally, some models rely on combinations of both sorts of abilities, such as the *structural alignment model* of Gentner and Wolff (1997), which proposes a stage where features of concepts are matched, and then one where a system of such features is formed by abstraction. (It is often suggested that this theory captures a great deal of what Black (1962) had in mind.)

Any of these models is based on our ability to work with concepts: to match their features and judge salience with respect to them, to recognize prototypes of them, or to abstract from them. Again, we see that such abilities are not language-specific, and are not part of linguistic competence. We see consistently in both philosophical and psycholog-

ical approaches to metaphor that the cognition of metaphor is part of our general ability to work with concepts, and not part of linguistic competence.

I shall label whatever is responsible for the cognition of metaphor the *metaphor system*. Our brief survey of ideas about the metaphor system in philosophy and psychology shows that it is safe to assume the metaphor system is distinct from linguistic competence. Beyond that, our survey leaves open just what the metaphor system is. It may be that the metaphor system is really just our most general cognitive system for working with concepts, or it may be a more specific part of that system. For our purposes here, we do not have to say any more.¹²

3.2. Assumptions about Language

I have concluded that the metaphor system is squarely in the conceptual realm, and distinct from linguistic competence. In doing so, I have already tacitly relied on some assumptions about language that I shall now make explicit.

I need to say enough about the cognition of language, specifically, about linguistic competence, to make clear that it is different from what we saw with the metaphor system. Our competence with language does not appear to be a matter of our general abilities to work with concepts: to sort similarities between concepts, to make abstractions from them, etc. Rather, linguistic competence is a highly specific ability, with its own specific organizing principles. (Indeed, it seems that we vary quite a bit in how well we comprehend metaphor, and generally how effectively we can work with concepts, while core language competence such as syntax shows remarkable uniformity in mature speakers.) Having this ability is typically described as the result of possessing a dedicated *language faculty* which determines the principles. This view is, of course, forcefully advocated by Chomsky (e.g. Chomsky, 1980, 1986b). Chomsky frequently describes a language faculty as a ‘mental organ’: a cognitive system that is special-purpose in much the same ways that organs like the heart and kidneys serve special purposes. Such a cognitive system is also frequently described as a ‘module’ in cognitive architecture, and indeed, the Chomskian view does posit many of the features that are characteristic of modularity in the sense of Fodor (1983). The language faculty appears to be highly domain specific, its operation is

mandatory, its workings are relatively inaccessible to consciousness, it shows characteristic patterns of development and breakdown, it has some features of informational encapsulation, etc. Even so, it is important to note that the Chomskian view does not make language an input module in Fodor's sense. (See Higginbotham (1987) for further discussion.)

I am generally sympathetic to a Chomskian view of language, and I shall often talk in Chomskian ways. However, to bear out the idea that metaphor and language cognition are different, we need only a few minimal features of the Chomskian view. So long as there is a distinct linguistic system, with principles not derived from abilities to work with concepts like abstraction and feature comparison, we have what we need. Assuming that linguistic competence is a matter of having a language faculty (in a mature state) guarantees this, but so would any other way of seeing distinct kinds of cognitive systems. If you want to insist, for instance, that the difference between linguistic competence and the metaphor system is a matter of learned abilities to apply more useful strategies in different kinds of settings, I need not argue against you here. (I would elsewhere!)

A number of other features of the Chomskian package are not at issue here, and can be ignored for our discussion. I shall say nothing about acquisition, either of language or of the ability to comprehend metaphor. I shall thus say nothing about questions of innateness. I shall also say nothing about linguistic universals and variation across languages, and issues of internalism and so-called 'I-language' versus 'E-language'.

With all these caveats about not needing the full force of the Chomskian view, I shall continue to talk about the language faculty (and I still think this is the best way to think about linguistic competence). There is one further issue about which I do want to be somewhat more careful. Even amongst people like me, who accept the general Chomskian view of the language faculty, there is lively debate about just what goes into it. As is common, I shall assume that syntax and phonology are within the language faculty. I shall also assume that in a very minimal way, some 'semantics' fall in the language faculty as well. More specifically, I shall assume that some structural representations relevant to meaning fall within the language faculty. In fact, these tend to look more like

syntax than like semantics once we find them (and hence, they are often described as part of the syntax-semantics interface). The sorts of things I shall assume to be within the language faculty are thus what we may loosely call 'grammar'. Importantly, though I shall talk about semantic properties within the language faculty, I am not assuming that these properties suffice to determine what is intuitively said by a sentence. Though (pace relevance theorists) I am inclined to think they do, it will not matter here. The only aspects of semantics we will talk about here are the ones that look like grammar.¹³

We have now identified two cognitive systems: the metaphor system is responsible for the comprehension of metaphors, while the language faculty is responsible for linguistic competence. I shall often talk about the latter as the *linguistic system*, to make it easier to talk about both at once, and to de-emphasize unneeded aspects of the Chomskian view. We have seen a few features of both systems, but what we really need is simply that they are distinct.

3.3. *The Linguistic–Metaphor Systems Interface*

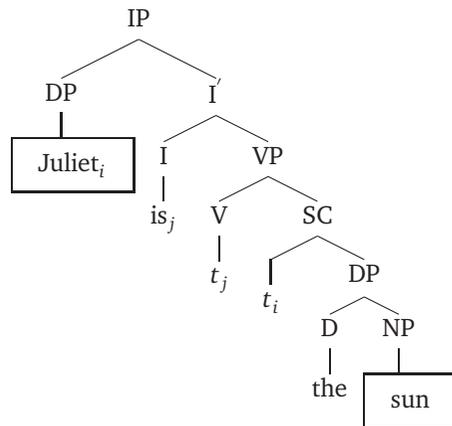
I have now reviewed why it is safe to assume that there are distinct cognitive systems at work in linguistic competence and metaphor. Once we have these systems, we know that they must interact, from the obvious fact that metaphors are expressed in language. Thus, there must be some way in which the two systems talk to each-other, or in the usual jargon, interface. I shall suggest that our constraint on metaphor from sections 1 and 2 is best seen as a facet of this interface.

I shall sketch, for discussion purposes, a simple model of how this interface might work. I shall focus on the issue of how a metaphorical interpretation gets assigned to a sentence, uttered in a context. The structure of the sentence, its sound, and the basic properties of the words in it, are determined by the linguistic system. This must then be operated on somehow by the metaphor system, to produce a metaphorical interpretation. For this to happen, the right information from the linguistic system must be passed to the metaphor system. The metaphor system must then act on whatever it takes in from the linguistic system, to produce a metaphorical interpretation of the sentence.

A natural and simple model of how this might happen is as follows. As the metaphor system is part of our ability to work with concepts,

only the conceptually rich elements of linguistic meaning suitable for abstraction, feature comparison, etc., are passed to the metaphor system. All other aspects of a sentence are stripped away in the interface.

To make this vivid, let us think about what the language faculty might produce. Here is a simplified version of one proposal for the syntax of an all-too-familiar sentence:



The details of this structure are not really important.¹⁴ All that matters is that linguistic theory tells us the structure of this sentence determined by the language faculty is very rich. In particular, it contains all sorts of things I have not explained, like subscripted t_i and t_j , $I\bar{P}$, I' and SC , which could not be operated on by the metaphor system. Our general system for manipulating concepts could not make sense of these technical language-specific notions.

We thus suppose that the metaphor system ignores all the structural parts of the sentence provided by the language faculty, and simply selects the concepts associated with the lexical items *Juliet* and *sun* (the boxed elements in the tree). Only these are passed to the metaphor system in the interface, leaving for the metaphor system simply:

(JULIET, SUN)

(I shall indicate concepts by SMALL CAPITALS.) The interface between the linguistic and metaphor systems, according to this idea, is one of generally erasing all linguistic structure provided by the linguistic system,

and selecting only concepts linked to lexical items for processing by the metaphor system.¹⁵

As this model of the linguistic–metaphor systems interface is the simplest reasonable one, let us call it the *simple model*. The simple model does have some plausibility. At least, it is quite plausible that lots of linguistic structure generated by the language faculty must be erased in the interface. The simple model simply takes the strongest available position on how much is erased.

The simple model faces many challenges, and ultimately, I shall argue it is wrong. But there are some shortcomings of the model which really show no more than that it is incomplete. I shall mention one that is particularly relevant to the issue we are examining here. The simple model is incomplete in failing to describe how complex metaphors can be part of larger sentences. Many of the metaphors we have already seen, for instance in (7) and (11), get their significance by building metaphorical interpretations into larger messages expressed in part by normal linguistic means. As we have seen, complex metaphors in sentences involving functional items like determiners are like this. At the very least, this means the simple model must be extended to provide for interaction in both directions across the linguistic–metaphor systems interface. Metaphorical interpretations must be fed back to the linguistic system for completion and packaging, at least.

In taking the strongest available position on what is erased in the linguistic–metaphor systems interface, the simple model captures the constraint on metaphor we found in sections 1 and 2. Along with everything else outside of concepts linked to lexical items, elements of functional categories are erased in the interface with the metaphor system, according to the simple model. The model thus predicts that functional categories cannot have metaphorical interpretations, as they are rendered invisible to the metaphor system by the interface.

The simple model is fairly modest in how much it explains about the constraint. It does not address why functional categories are erased in the interface, even though some of them—like determiners—have the kinds of contents which can have metaphorical interpretations when expressed by lexical categories.¹⁶ But even so, it does give us a model of how the constraint works: it works as a constraint on the linguistic–metaphor systems interface, rather than, say, as a constraint on metaphor

processing itself. It also offers an answer to the question of how far the constraint really goes. It implies that anything outside of the concepts linked to lexical items, and so anything that could count as linguistic structure, will fail to have metaphorical interpretations. More strongly, any such structural elements will be invisible to the metaphor system.

Though I think an interface model is the right way to model the constraint, I shall argue in the following section that this latter prediction is incorrect. I shall show that there are aspects of linguistic structure, importantly similar to functional categories in some respects but appearing within the lexicon, which must be available to the metaphor system. Thus, the simple ‘erase practically everything’ model must be replaced with a much more selective model of what gets passed to the metaphor system in the interface.

4. METAPHOR AND LEXICAL STRUCTURE

We have now seen that functional categories cannot be interpreted metaphorically. This shows us something about functional categories, and it provides a constraint on metaphor. In the last section, I suggested this constraint is best understood as a constraint on the linguistic–metaphor systems interface, and I sketched a simple model of the interface which accounts for the constraint.

The main goal of this section is to argue that the simple model is wrong. In particular, its prediction that no aspects of linguistic structure are passed from the linguistic system to the metaphor system is incorrect. We already have seen reasons to think the simple model is inadequate, and at least needs to be extended, but this will show that the simple model is really on the wrong track. It will show that the model of erasing all aspects of linguistic structure at the linguistic–metaphor systems interface must be replaced by one where some but not all aspects of linguistic structure are selected by the metaphor system. Showing this will help us to better understand the extent of the constraint on metaphor we have uncovered.

I shall substantiate this claim by showing that some structural elements related to lexical meaning must be available to the metaphor system. The main argument about metaphor will be given in section 4.4. Before this can be done, however, I shall have to present some

ideas about linguistic structure in the lexicon in sections 4.1 and 4.2, and then revisit the simple model in light of those ideas in section 4.3.

4.1. Linguistic Structure in Lexical Meaning

Some aspects of the meanings of words have something to do with linguistic structure—with grammar—while others do not. If you want to know the difference in meaning between *fear* and *frighten*, a linguist will have something interesting to tell you, and what they say will be mostly about structure, indeed about grammar. In contrast, if you want to know the difference in meaning between *red* and *blue*, you will have better luck asking a psychologist.

There is not anything close to universal agreement over what makes certain things fall within the purview of linguistic analysis. But here is the sort of picture with which I shall work. The picture supposes that lexical meaning is a combination of two sorts of factors. One is linguistic structure, which can be mapped and explained by a good linguistic theory. The other is more idiosyncratic and conceptual, and not the sort of thing linguistic theory proper is likely to shed light on. The main difference between *fear* and *frighten* is structural: it has to do with which argument is the *experiencer*. The main difference between *red* and *blue* is conceptual: it has to do with whatever goes into our grasp of colors. We label this *idiosyncratic* from the point of view of linguistic theory (not from psychology in general), as it is not something we can organize or explain within linguistics itself.¹⁷

The structural side of lexical meaning has a lot in common with functional categories. It consists of a small number of elements, determined by the linguistic system. They have contents, but like functional categories, these contents tend to be abstract in nature, and typically need to be completed by idiosyncratic features of lexical meaning to get us what we think of as the meaning of a word. We have already found tempting the assumption that this sort of content is home-grown in the linguistic system (though this is not essential to the arguments presented here). In contrast, idiosyncratic content represents established points of interface between the lexicon and general conceptual systems. Again, we may remain neutral on just how that interface works, though it is tempting to think of the lexicon as containing pointers to the relevant concepts. I shall continue to use the term *idiosyncratic* to describe

the non-structural or conceptual side of lexical meaning, with the reminder that it means only idiosyncratic with respect to linguistic theory. (Ultimately I shall argue that both structural and idiosyncratic contents can be inputs to the metaphor system, so it is best not to insist on only one being conceptual.)

I shall work with a specific example of this sort of approach to lexical meaning: the theory of *aspectual classes* of verbs. This theory has a fine pedigree, stemming from ideas of Aristotle, through work of Ryle (1949) and Vendler (1967), to the extensive development by Dowty (1979), and many others. This theory, like pretty much everything in lexical semantics, is highly controversial. But it is an established approach to some important aspects of the lexical semantics of verbs, and it is sufficiently well-known to provide a familiar example. Also, as we will see, it helps us to focus on some important data we can apply to questions about metaphor.¹⁸

4.2. A Brief Tutorial on Aspectual Classes

Before getting to the issues of metaphor that are our main concern, I shall quickly review some of the basics of the theory of aspectual classes. I shall be very brief, and gloss over a number of disagreements about how the theory should be formulated, and criticism of it. My presentation will fairly closely follow Rothstein (2004). Other extended discussions include Levin and Rappaport Hovav (2005), Smith (1997), Van Valin (2005), and Verkuyl (1993), as well as the classic Dowty (1979). See any of these for references to the extensive literature.

The theory of aspectual classes is a theory of verbs. It groups verbs into four classes.

States: *love, know, have,*

Activities: *run, walk, swim,*

Achievements: *die, recognize,*

Accomplishments: *paint a picture, draw a circle,*

Different versions of the theory propose slightly different classifications, but this traditional one will suffice for our purposes. It has also been a

matter of intense discussion whether these are properly classes of verbs or of verb phrases. I shall ignore this issue, and shall continue to talk about classes of verbs. (See the references above for discussion.)

Verbs generally describe events.¹⁹ When they do so, they describe events in certain ways, with certain structural features. The aspectual classes indicate at least some features of how verbs depict events, and are a clue to the kinds of structural components of verb meaning that go into building such depictions. They are thus a clue to the linguistic structure that might be found in the lexicon.

On one way of looking at them, the aspectual classes point to two features of how events may be depicted. One is usually called *telicity*. Telic verbs describe events as having an endpoint or telos, which is a kind of culmination point for the event described. Achievements and accomplishments are telic: describing an event as a dying or a painting of a picture includes describing it as as having a point at which it is finished. States and activities are not telic: describing an event as a loving or a running does not include describing it as having such a culmination point.²⁰

The other feature is what Rothstein (following work of Landman, 1992) calls *stages*. Describing an event with an achievement verb describes it as not having any temporally extended stages in which the depiction is occurring, as it describes the event as virtually instantaneous. The same goes for states. Though they may be extended in time, they need not be, and do not have stages in which the state is occurring. They do not, in other words, have distinguishable parts where a process is ongoing. It is sometimes said that they are thus non-dynamic. In contrast, both activities and accomplishments describe events as having stages where the processes described are occurring. Any event of running or painting a picture has stages where running or painting is happening.²¹

The terms *telic* and *stages* are labels for sorts of event depictions. I shall not go into the details of how they may be analyzed.²² There are some typical linguistic reflexes of these properties, which serve both as tests for which aspectual class a verb falls under, and further guides to the nature of the properties. One of the key tests for telicity is occurring in constructions with *for/in x time*: telic VPs typically occur with *in x time* while atelic ones occur with *for x time*. One of the key tests for

stages is occurring (easily) in the progressive: states and most achievements cannot occur in the progressive, while activities and accomplishments can. The literature includes a number of other tests (many from Dowty, 1979). For instance, the so-called ‘imperfective paradox’ can help distinguish activities from accomplishments: for activities, we typically have the progressive entailing the past, while this is not typically so for accomplishments.

All these tests have difficulties. Dowty (1979) raised questions about the imperfective paradox test. There is a well-known class of achievements which do appear in the progressive: so-called *progressive achievements* such as *The old man is dying*. (We will make use of progressive achievements below.) Activities can occur with *in x time* under ‘time of onset’ readings, such as *Mary will swim in ten minutes*. Hence, the tests must be applied with care (and hence, some controversy about how solid classification by aspectual classes is, and about what it classifies). Even so, we have enough to at least get an inkling of how different verbs can differ with respect to structural aspects of their meanings.

If aspectual classes are a clue to structural aspects of verb meaning, how do structural and idiosyncratic features of a verb’s meaning determine its aspectual class? Dowty (1979) originally suggested that aspectual class is determined by a small group of structural operators that apply to idiosyncratic contents. The operators include *DO* and *BECOME*. This view has been modified by Rothstein (2004) to fit into a framework where verbs are treated as predicates of events, and she adds an event predicate *Cul* and an event summing operator \sqcup . (There are some technical issues about just what sort of summing operator is required, but they will not be relevant here.) Using these, the Dowty-Rothstein proposal is that the aspectual classes correspond to the following frames for verb meanings.

- States: [–stages, –telic].
Frame: $\lambda e.P(e)$.
- Activities: [+stages, –telic].
Frame: $\lambda e.(DO(P))(e)$.
- Achievements: [–stages, +telic].
Frame: $\lambda e.(BECOME(P))(e)$.

- Accomplishments: [+stages, +telic].

Frame: $\lambda e.\exists e_1\exists e_2(e = e_1 \sqcup e_2 \wedge (DO(P))(e_1) \wedge Cul(e) = e_2)$.

P is an idiosyncratic content (or a pointer to one). According to this proposal, the idiosyncratic meanings of verbs are basically stative predicates of events. Hence, the frame for state verbs has no additional structure. Each of the other aspectual classes adds linguistic structure, to provide for the features of the relevant aspectual class. This proposal thus holds that accomplishments have the most linguistic structure, and states the least.²³

As given, this is more of a template for a proposal than a full proposal itself. An analysis of the operators *DO* and *BECOME*, and an explanation of how they generate the aspectual classes and account for their features, is needed to fill in the template. However, I shall assume that the basic idea behind the operators is clear enough, and that they intuitively divide verbs along the lines the aspectual classes indicate. I shall thus assume we may work with these frames as a model of how idiosyncratic and structural elements may enter into a verb’s meaning. I do so guardedly, as we have already seen how controversial a proposal like this is bound to be.

This theory of verb meaning gives at least a little substance to the picture of lexical meaning discussed in section 4.1. It locates an idiosyncratic part of a verb’s meaning, indicated by *P*, which is a point of interface with the general conceptual system. It also locates structural elements which provide a frame for the verb’s meaning. These are drawn from a limited stock, and appear in ways that are determined by the linguistic system. As we have come to expect from such structural elements, they have meanings that are highly abstract, and have as much to do with grammar as with what we intuitively think of as meaning.²⁴

4.3. The Simple Model and Lexical Meaning

Now that we have at least a rough idea of how structural and idiosyncratic elements might appear in the lexicon, we need to return to the simple model of the linguistic–metaphor systems interface. The simple model proposes to erase anything that could count as linguistic structure, including such functional categories as determiners that do have

some contents associated with them. Anything distinctive of the linguistic system is stripped away in the interface, according to this model. In describing the simple model in section 3.3, I said that it simply selects the concepts associated with lexical items (of genuine lexical categories!). Now that we have seen that within the lexicon, we also have structural elements, we need to refine our statement of the simple model to take these elements into account.

If the simple model is to erase any structural elements provided by the linguistic system, it must do more than simply target lexical items. It must target the parts of lexical entries that provide idiosyncratic content, and erase the structural frames in which these idiosyncratic contents appear. The simple model should be refined to pass to the metaphor system only those elements within the lexicon we represented above by *P*.

To take an example, the lexical entry for the achievement verb *die* is $(BECOME_{(DEAD)})(e)$, where *DEAD* is the idiosyncratic content of the state of being dead. The linguistic system depicts this content with the verb *die* as an achievement, rather than, say, a state. According to the simple model, all that is passed to the metaphor system from *die* is the idiosyncratic content *DEAD*. Thus, the interface between the lexicon and the metaphor system, for this case, looks something like:

$$\begin{array}{c} (BECOME_{(DEAD)})(e) \\ \Downarrow \\ DEAD \end{array}$$

On this model, everything but the idiosyncratic content is erased in the interface with the metaphor system.

The simple model, thus refined, is an appealing picture. As we discussed in section 3.3, it reflects the idea that the metaphor system is part of our broad ability to work with concepts, and so only conceptually rich elements should be passed to it by the linguistic system. Thus it proposes to erase all structural elements, including those within the lexicon. As I also mentioned, the simple model is very aggressive in what it erases. It erases structural elements that have genuine contents, including functional categories and elements like *BECOME*. If we continue to speculate that these elements have contents that are home-grown by the

linguistic system, this aggressive model might seem natural. It erases anything distinctive of the linguistic system, and passes to the metaphor system only those elements that are naturally part of the broader conceptual system to begin with. It essentially proposes that the interface to the metaphor system is exactly the interface from the lexicon to the conceptual system.

Though this is very natural, I shall now argue it is wrong. I already mentioned that there are many reasons to doubt the simple model. But the ones I shall offer in the next section show it to be on the wrong track.

4.4. Metaphor inside the Lexicon

We may now finally get back to metaphor. To show the simple model is wrong, I shall argue that metaphorical interpretations of verbs can be sensitive to aspectual structure, and thus, the linguistic–metaphor systems interface must pass information about the structural frames of verbs we reviewed in section 4.2 to the metaphor system.

Unfortunately, my arguments for this conclusion must be somewhat indirect. Unlike the case of functional categories we explored in sections 1 and 2, it is impossible to simply point to some element of aspectual structure and ask if it is interpreted metaphorically. We cannot do this, as aspectual structure is buried within the lexicon, and is not visible in such a way that we can directly localize metaphorical interpretation to it. Instead, I shall argue that we can see properties of metaphorical content that reflect aspectual structure, and in particular that we can isolate differences in aspectual structure that lead to differences in metaphorical content. This, I shall argue, is best explained by a model which passes elements of aspectual structure to the metaphor system.

To do this, I shall examine a couple of verbs in detail, looking both at their aspectual structures, and the metaphorical interpretations they receive. I shall begin with the verb *blush*, which has received substantial discussion in the literature. It is an especially useful case, as it allows for a cross-linguistic comparison that will help to isolate aspectual from idiosyncratic content.

Blush is an activity verb. We can run a few of the standard tests to show this.

- For/in:

- (12) a. Mary blushed for an hour.
b. * Mary blushed in an hour.

(Indicates –telicity.)

- Occurs (easily) in progressive:

- (13) Mary is blushing.

(Roughly indicates +stages.)

- Progressive entails past:

- (14) Mary is blushing ENTAILS Mary blushed.

(Distinguishes activities from accomplishments. Further indicator of telicity or related properties.)

We thus know a little about what the internal structure of *blush* is like. It has the frame of an activity verb, and so has a lexical entry like $(DO(\text{RED CHEEKS}))(e)$, where RED CHEEKS is the idiosyncratic content that goes with *blush*. The English verb *blush* describes an activity of DOing (or being) red in the cheeks.

There are lots of metaphors with *blush*. Here are a few, many drawn from Shakespeare, as usual:

- (15) a. The sun of heaven methought was loath to set, But stay'd and made the western welkin blush (*King John* VV).
b. Let my tears stanch the earth's dry appetite; My sons' sweet blood will make it shame and blush (*Titus Andronicus* III.I).
c. The windows blush with fresh bouquets (Oliver Wendell Holmes, *The Autocrat of the Breakfast-Table*).

These are all metaphors of personification (cf. Lakoff and Johnson, 1980). As such, they have more content than simply saying that the object in question is red. (Note: *welkin* means sky. The *Titus Andronicus* quote is part of a complex extended metaphor of personification of the earth that runs throughout the whole scene.)

One of the reasons for focusing on the verb *blush* is a much-discussed cross-linguistic comparison that shows structure in the lexicon at work. The Italian translation of *blush*, the verb *arrossire*, is a faithful translation of idiosyncratic content, but is an achievement rather than an activity verb.

Conceptually, *blush* and *arrossire* have the same (idiosyncratic) content. They both describe the reddening of the cheeks. For both English and Italian speakers, this is the result of a bodily process, it is associated with embarrassment, and induces the same patterns of responses. As far as idiosyncratic content goes, the two verbs really are the same.

Even so, *arrossire* is an achievement verb (Levin and Rappaport Hovav, 1995; McClure, 1990). Here are some tests.

- For/in:

- (16) a. G è arrossito in un secondo.
G has blushed in one second.
b. * G è arrossito per 10 minuti.
G has blushed for 10 minutes.

- Progressive does not entail past:

- (17) G sta arrossendo DOES NOT ENTAIL G è arrossito.
G is blushing DOES NOT ENTAIL G has blushed.

(Data from McClure (1990) table 4.) Thus, *arrossire* has a different frame from *blush*, even though it has the same idiosyncratic content. We have:

- (18) a. Blush: $(DO(\text{RED CHEEKS}))(e)$.
b. Arrossire: $(BECOME(\text{RED CHEEKS}))(e)$.

Even without a full semantics for the *DO* and *BECOME* operators, this makes vivid what the structural difference is. Italian represents the event as one of becoming red in the cheeks, whereas English represents it as a doing of being red in the cheeks. Different languages can make different choices about how to depict the same idiosyncratic content, by putting it in different frames.

Now we have two verbs, *blush* and *arrossire*, which have the same idiosyncratic content but different structural frames. The simple model, which simply selects idiosyncratic content, predicts that we should see the same metaphors with the two. But this is not so. Rather, we find that metaphors with the two verbs typically reflect their aspectual structures. Compare:

- (19) a. The sky blushed.
b. Il cielo arrossí.

These express metaphors of personification. The simple model says they should have the same metaphorical content, but speakers report that they do not, as they report that the aspectual structure of an activity is part of the metaphorical content in (19a), while the aspectual structure of an achievement is part of the metaphor in (19b).

Such judgments, however, do not tell us quite enough. I have found that, especially if a target sentence is surrounded by other text, informants are sometimes willing to import event structure into a metaphor from clues in the surrounding text, rather than from the target sentence itself.²⁵ So, we would like to sharpen our evidence to show that it is the aspectual structure of the verbs that is responsible for the difference in metaphorical content. This will help show that aspectual structure is passed from the linguistic to the metaphor system.

One further test we can run to do this involves putting the verbs in sentences that highlight, or even require, specific aspectual structure, such as the kinds of sentences that we use to test for aspectual structure in the first place. If the simple model were right, and metaphor interpretation ignored aspectual structure, then we might expect metaphorical interpretations to be available for sentences that place a verb in an environment that typically does not allow its aspectual structure.

This does not happen. Consider:

- (20) a. Mary blushed all day long.
b. * Mary arrossí tutto il giorno.

These are word-for-word translations, but the Italian version is judged unacceptable. The temporal modifier is acceptable with an activity but not an achievement.

Metaphor of personification does not change this:

- (21) a. The sky blushed all day long.
b. * Il cielo arrossí tutto il giorno

If (21b) were acceptable, it would support the simple model. But its failing to be acceptable does not yet show the simple model is wrong. It does not, as its unacceptability may simply be an effect of grammar, independent of metaphor interpretation. It may be, for instance, that metaphor interpretation works the same in both the English and Italian cases, but then the linguistic system takes over to build a sentence around a metaphorical interpretation, and that rules out (21b) as ungrammatical.

To try to control for this, we need to try to correct the sentence for grammar, with as little change as possible. Asking an informant to replace (21b) with an acceptable sentence gives:

- (22) Il cielo restó arrossito tutto il giorno.
The sky stayed blushed all day long.

My informant reports this is an acceptable sentence.

Even so, my informant reports that (22) is only marginally intelligible for meaning, in spite of being a metaphor. Yet if metaphor interpretation ignored aspectual structure, metaphorical interpretations of (22) should be readily available. At the same time, the marginal intelligibility is nicely explained by supposing that the metaphorical interpretation interprets the verb *arrossire* with its achievement structure. We thus have some evidence that metaphor interpretation not only reflects aspectual structure, it is sensitive to the aspectual structure of the verb it interprets. We have evidence that aspectual structure is passed from the linguistic system to the metaphor system.

We get similar results when we look at cases where we can force a phrase to get interpreted as within two different aspectual classes. One example is that of progressive achievements: cases where an expression that is normally an achievement is able to appear in the progressive, and the whole progressive construction is interpreted as an accomplishment. (Recall, usually achievements cannot occur in the progressive at all.) For instance:

- (23) a. John died.
b. John was dying.

In (23a) we have an achievement, while in (23b) we have an accomplishment.²⁶

The idiosyncratic content of the verb is the same in both cases, but even so, there is a difference in meaning corresponding to the different aspectual classes. This difference is clearly preserved in metaphors of personification, as we see with:

- (24) a. My idea died.
b. My idea was dying.

Informants give me consistent judgments of difference in meaning between these two (regardless of whether or not I also show them instances of metaphorical versus non-metaphorical interpretation, or of non-metaphorical progressive achievements).

The simple model predicts we should have the same metaphorical contents available for (24a) and (24b), as both have the same idiosyncratic content. Thus, we again have evidence that the simple model is wrong. The two sentences in (24) do have different inflectional structure, as one is in the progressive and the other is not. But this is more linguistic structure, of just the sort the simple model erases. So, as far the simple model goes, these sentences should pass exactly the same content to the metaphor system, and thus should receive the same metaphorical interpretations. This is not what happens, so the simple model must be wrong.

The evidence we have seen from *arrossire* versus *blush*, and from progressive achievements, gives us good reasons to think that metaphorical interpretation is sensitive to aspectual structure as well as idiosyncratic content. Thus, the evidence gives us good reasons to think that the simple model is incorrect. However, an alternative explanation of the data we have just seen remains available to the simple model, which exploits the potential extension of the model I mentioned in section 3.3. It might be that the data we have seen is not the result of aspectual structure being passed from the linguistic system to the metaphor system. Rather, it might be that the metaphor system operates only on idiosyncratic content, as the simple model says, but it then passes the metaphorical interpretations of idiosyncratic content back to the metaphor system, where it is placed back in the same structural frames it was taken out of. Thus, the whole metaphorical sentence might reflect aspectual structure, but because of effects generated on the linguistic

side rather than the metaphorical side of the interface. This is not exactly what the simple model proposes, but it extends the simple model in a way that does not undermine its basic point.

I shall now argue against this alternative explanation directly. To do so, I need to sharpen our results still further. I need to show not only that we find our understanding of a whole metaphorical sentence to be sensitive to aspectual structure, but that there are features of the metaphorical interpretation of particular expressions which would not make sense unless the metaphor system had access to aspectual structure. Thus, I shall sharpen our results to show that the metaphor system must see operators like *DO* and *BECOME*. The strongest claim one could make here is that these operators themselves can receive metaphorical interpretations. I shall argue directly only for the weaker claim that the metaphor system must have access to these operators. I am inclined to think the stronger claim is true, and the evidence I shall provide does make it seem plausible, but the evidence I have so far is not quite sufficient to really establish it.

The argument I shall advance focuses on an indirect effect on interpretation with the modifier *when I looked at him/her/it* in its punctually locating sense. This modifier applies to achievements and to states, but not to accomplishments and activities (excluding the onset of activity reading). When applied to achievements, it typically triggers an implicature of causal agency on the part of the speaker. For instance:

- (25) a. John died when I looked at him.
b. Mary arrossí quando la guardai.
Mary blushed when I looked at her.

Though the verbs in question are not themselves causal, all of these indicate that the speaker had some causal effect relating to the outcome. (In the Italian case, I am told, this implication is very clear). This effect is generally absent for states:

- (26) a. John was dead when I looked at him.
b. John was happy when I looked at him.
c. John knew me when I looked at him.

We do not typically see the implicature in any of these. (We might generate it with enough additional contextual information, of course.)

I have labeled this effect an implicature, and it shows the signs of being a conversational implicature. It is, for instance, cancellable in context. It is perfectly coherent to say *John died when I looked at him, but it was not my fault*.

The ready availability of the implicature for achievements but not states suggests that a key trigger of the implicature is aspectual structure. Here is one rough story about how such an implicature might arise. The frame for achievements, but not states, includes a *BECOME* operator. This contributes content to achievement sentences. Though it is highly abstract and structural in nature, it is enough for us to assume that normally, something *BECOMES* only if it is caused to become. We also, in cases like (25), have enough information to tend to infer that it is the speaker who did the causing. We hence wind up with the implicature of the speaker causing the outcome in (25). In contrast, states have no operators in their frames, and in particular no *BECOME* operator, so we do not draw this implication in (26). The precise details of how the implicature is calculated will not really matter. All that is important is that it relies on aspectual structure.²⁷

Given the kinds of data we have already seen, it should come as no surprise that this effect is preserved in metaphorical interpretation. We see the implicatures with metaphorical interpretations of achievements, as in:

- (27) a. My idea died when I looked at it.
 b. Il cielo arrossí quando lo guardai.
 The sky blushed when I looked at it.

Likewise we do not see it with metaphorical interpretations of states, as in:

- (28) a. My idea was dead when I looked at it.
 b. My idea was happy when I looked at it.
 c. My idea knew me when I looked at it.

Regardless of whether the interpretation is metaphorical or not, we get the same pattern of typical implicatures.

So far, this is effectively more evidence along the lines we have already seen. But if we pursue the implicature of speaker causation in

metaphor a little further, it will help us to identify an important aspect of metaphorical interpretation.

To do so, we need one more independent piece of the puzzle: an observation about how causal notions can get interpreted metaphorically. In some cases causal expressions can themselves be interpreted metaphorically, even if they have extremely thin or ‘light’ meanings. We see this with *make* in:

- (29) The sun of heaven methought was loath to set, But stay'd and made the western welkin blush (*King John* VV, repeated from 15).

In this example, *make* is part of a complex metaphor of personification (and a rather elegant metaphor, often missed in a not-much-loved play). Obviously *blush* is interpreted metaphorically to produce this complex metaphor; but furthermore, I shall argue, *make* is interpreted metaphorically as well.

The metaphorical interpretation of *make* is somewhat difficult to detect, as the literal meaning of *make* is already very broad and ‘thin’, and can apply quite liberally to all sorts of situations. Taken literally, *make* can describe all sorts of relations between celestial bodies. For instance, it appears to be literal in:

- (30) This kind of cloud makes the Moon turn red.²⁸

In our extended metaphor (29), we see something different. This metaphor describes an event of making in a personified sense: making in the sense of what a guest staying too long does when he makes his host blush. This is not merely causation, as it includes content specific to human agency and human responses. This goes beyond the literal meaning of *make*, which is a clear sign that *make* is interpreted metaphorically as part of the metaphor of the sun and sky personified.

For more evidence that we can give *make* a metaphorical interpretation involving agency, consider:

- (31) The huge cardinal made my head hurt.

This sentence has a non-metaphorical reading. Depending on the context, it might be understood as conveying something like *Thinking about the huge cardinal made my head hurt*. But it also has a metaphorical interpretation, where it is the cardinal (number!) itself that makes the

speaker's head hurt. This is a very natural reading of the sentence. In it, we see a personification not just of the cardinal number, but of the action the cardinal takes. This again involves a metaphorical interpretation of *make* involving human agency.

The availability of multiple readings for (31) suggests another test:

- (32) */?? The sun made the sky blush and the temperature rise by 10 degrees.

This sentence is bad, or marginal at best, which is evidence that the overt and elided occurrences of *make* receive different readings. The first overt one receives a metaphorical reading of human agency, while the second elided one does not.²⁹ We may conclude that *make* can be interpreted metaphorically in the course of a metaphor of personification, in spite of having a highly general and abstract literal meaning.

Now, we are ready to put the pieces together and make the main argument. From the observations about implicatures of causation and about metaphorical interpretations of causal notions, we can argue that *BECOME* must be available in metaphor interpretation. First, we have seen that the implicature of the speaker causing the outcome in the construction in (25) requires the *BECOME* operator, and is typically not present without it. (More generally, the implicature requires the right aspectual structure from the verb.) This implicature can occur with either metaphorical or non-metaphorical interpretations. Furthermore, from our observation of metaphorical interpretations of causal notions, we may conclude that in metaphors of personification, the implicated causal relation between the speaker and the outcome may also be interpreted metaphorically. In (27), the implicated causal relation receives the same kind of metaphorical interpretation that *make* does in (29). The implicature in these cases is not merely that the speaker enters into some causal relation leading to the outcome. As part of the personification metaphor, the implicature is enriched to include that the relation is the kind that occurs when one person induces another to blush. The causal relation is itself interpreted as personified, just as we saw with *make*.

This shows that the metaphor system is active in generating the implicature of causation. As this implicature is generated in part by the presence of the *BECOME* operator, we may conclude that the metaphor system acts on inputs including this operator. Thus, the linguistic-

metaphor systems interface must make the *BECOME* operator available to the metaphor system. The metaphor system sees *BECOME* (and perhaps other elements of aspectual structure). This is enough to show that some elements of linguistic structure are passed to the metaphor system. Hence, the simple model, which proposed that all elements of linguistic structure are erased in the interface, must be wrong.

I believe it is plausible that the *BECOME* operator itself receives a metaphorical interpretation in metaphors of personification, with the same enriched content as the causal notions receive. It is, I believe, interpreted as becoming in the sense that people become flushed, rather than the abstract sense of becoming that is the operator's literal content. I have not explicitly argued for this stronger thesis, but instead for the weaker thesis that the metaphor system must in some way have *BECOME* available. The weaker thesis is slightly easier to substantiate, as the implicature of causation has a much more specific content than the highly abstract *BECOME* operator, which makes it easier to provide evidence that the implicature may be interpreted metaphorically. It might be possible to find related arguments for the metaphorical interpretation of other elements of lexical structure. For instance, if we accept analyses of causal verbs which include a structural element *CAUSE*, then we might run similar arguments to show that this element can be interpreted metaphorically. I shall leave this task to another occasion.

I shall conclude this section with one final conjecture about how the metaphor system can make use of aspectual structure. Metaphor can *add* aspectual structure. Consider:

- (33) The patient is blue.

Metaphorical interpretation: the patient died.

Assuming that this is indeed a metaphorical interpretation, it represents the event mentioned in (33) as an achievement, even though *being blue* is a state.³⁰ The metaphor system is thus able to add its own aspectual properties. This does not tell us anything directly about how the metaphor and linguistic systems interface, but it does show that the metaphor system itself must have concepts available for aspectual properties. This does not by itself guarantee that the interface will pass aspectual elements to the metaphor system. We have already seen that functional categories are not passed, even though the metaphor system has avail-

able concepts corresponding to their contents. But it at least suggests that the metaphor system can understand aspectual elements, and so it can receive them from the linguistic system.

I have now made my case that the metaphor system must see aspectual structure. I have provided evidence that aspectual structure affects metaphorical interpretation, and that specific metaphorical interpretations require the metaphor system to have aspectual structure available. Though I grant that this evidence is somewhat indirect, I believe its most plausible explanation is that aspectual structure can be passed to the metaphor system in the linguistic–metaphor systems interface. Thus, some linguistic structure, within the lexicon, is visible to the metaphor system.

This shows that the simple model, which supposed that all elements of linguistic structure are erased in the interface, is wrong. Some elements of linguistic structure, including functional categories, are erased; but some, including aspectual structure in the lexicon, are not.³¹

5. RESULTS AND SPECULATIONS

We can break the results of this paper into three groups: results about what kinds of elements get interpreted metaphorically, results about the notion of functional category, and results about the interface between the cognitive systems responsible for metaphor and language.

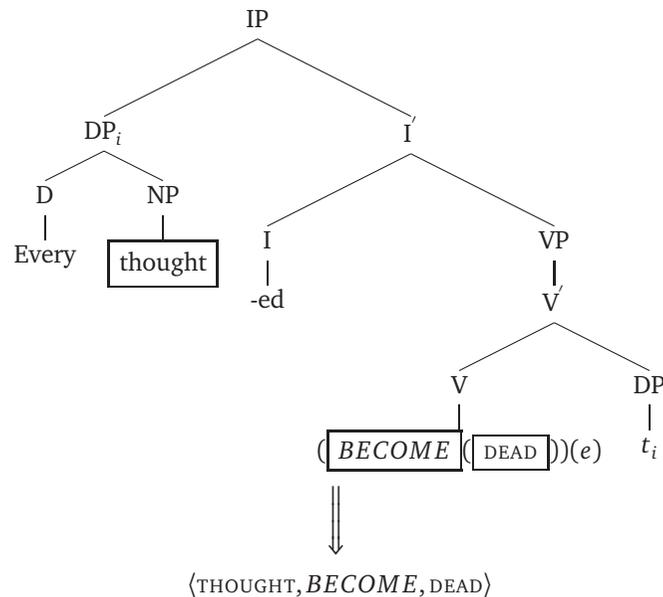
First, what gets interpreted metaphorically. We have seen that though expressions of the major lexical categories can be interpreted metaphorically, determiners, tenses and other functional categories cannot. This is so, even though lexical expressions with substantially the same contents as functional expressions can be interpreted metaphorically. We have also seen some reasons to suspect that elements of aspectual structure within the lexicon might also be capable of metaphorical interpretation, but the evidence I have presented here does not fully decide this.

Second, the notion of functional category. The correlation between functional categories and lack of metaphorical interpretation shows that we can use the availability of metaphorical interpretation as a probe for functional versus lexical status. At the same time, the availability of metaphorical interpretation gives some substance to the idea that

functional categories have only thin meanings. However, the metaphor probe works only at the categorical level, to distinguish functional from lexical categories. It does not offer an across-the-board test for structural versus non-structural status, as the possibility remains open that functional structure within the lexicon might be interpreted metaphorically. Even so, the notion of functional category is an important one in its own right, and metaphor can be a tool for exploring it.

Finally, the interface. The fact that expressions of functional categories cannot be interpreted metaphorically, even though correlate terms of lexical categories can, is a substantial constraint on metaphor interpretation. I argued that it is best understood as a constraint on the interface between the cognitive systems responsible for linguistic competence and metaphor comprehension. Functional categories cannot be interpreted metaphorically because they are not passed to the metaphor system in the interface. I went on to explore how extensive this constraint is. I considered a simple model of the linguistic–metaphor systems interface, which proposes that all linguistic structure is erased in the interface. I argued this simple model is wrong, as some features of linguistic structure within the lexicon must be visible to the metaphor system.

What should replace the simple model? We cannot at this point say. Even so, we can conclude that the interface between the linguistic and metaphor systems is highly selective about which structural elements it passes to the metaphor system, and which ones it does not. It selectively erases functional categories (categorical linguistic structure) but passes aspectual structure (lexical linguistic structure) to the metaphor system. At the very least, the interface must look something like:



Again, the syntactic analysis of the sentence has been simplified. The important point is that the metaphor system somehow targets lexical elements, including structural and idiosyncratic content in the lexicon, and ignores functional elements. The simple model offered an elegant and natural story about functional categories, but we have seen that at least this more complex *selective* model is required.

I do not have a good explanation for why the linguistic–metaphor systems interface works this way. It is a topic for further investigation. However, one speculative point is worth mentioning. Somehow, it appears that components of lexical meaning, even structural components, have some feature that makes them visible to the metaphor system when functional categories are not. There is already a suggestion in the literature which might indicate what that feature is. Higginbotham (1985, 2000) and others have proposed that all lexical categories include an *event* argument, while functional categories do not.³² It is easy to speculate that the metaphor system might be sensitive to whether or not something is involved in describing eventualities. If the linguistic system encodes this feature of describing eventualities in terms of having an event argument (or perhaps related properties like assigning or

discharging theta-roles), then it might offer a way to better explain the constraint on the interface between the linguistic and metaphor systems I have proposed here.³³

Speculation aside, we do have a non-trivial constraint on how the metaphor system and the linguistic system interact.

Versions of this paper were presented at A Figure of Speech: Conference on Metaphor, University of Latvia, Riga, 2007; the UC Davis ‘Salon’, 2008; and the Paris-Stanford Conference on Meaning, Context and Communication, Stanford University, 2008. I am grateful to all the participants at those events for valuable discussions. Thanks also to Robert May, Adam Sennet, Paul Teller, and Michael Wedin for discussions and comments on earlier drafts, and to Greg Damico and Francesca Merlin for judgments and help with Italian. The students in my classes at UC Davis in Fall 2008 also provided judgments. Finally, special thanks to Liz Camp for answering my many naive questions about metaphor.

Notes

¹ This point is made elegantly by Stern (2000).

² We can also dispense with any strong claim about literal meaning here. What is important is that we can recognize a contrast between more or less literal and metaphorical, and localize it to parts of sentences. I am a firm believer in literal meaning, and shall talk about it throughout this paper. But so long as we suppose that we can recognize some cases of metaphor, weaker assumptions about the literal would suffice.

³ There is a substantial debate over whether every determiner expresses a quantifier, or if there is a distinguished subclass of quantifying determiners. There are also non-determiner quantifiers like adverbs of quantification. I shall ignore these issues here.

⁴ In traditional lists of figures, the use of a term outside of its category is often labeled *antimeria*.

⁵ I have relied on the Opensource Shakespeare concordance of Shakespeare’s complete works, at <http://www.opensourceshakespeare.org/concordance/>. I should mention that Shakespeare’s English is in some important linguistic ways different from ours. For instance, we find such constructions as:

(i) Two glasses where herself herself beheld (*Venus and Adonis* 1151).

This is not modern English. Regardless, the determiner system in Shakespeare’s English seems to be close enough to modern English to make Shakespeare’s metaphor-rich texts a good place to explore what metaphors are available.

⁶ You can see this in any recent syntax text, or in the important discussion of Grimshaw (2005a). The leading idea is that both noun and verb phrases live inside larger syntactic structures built up from functional elements. Where we would have intuitively looked for a sentence, for instance, this view holds that we have an inflection or tense phrase: a phrase built around the functional category of inflection or tense. Lexical elements fall

under these superstructures. Particularly, verb phrases fall under inflection phrases, while noun phrases fall under determiner phrases. See Fukui (2001) for an overview.

⁷ I have so far said nothing about the status of prepositions. It is a matter of some dispute in the syntax literature whether they should be counted as lexical or functional. Since Chomsky (1970), it has been common to see the lexical categories as themselves defined in terms of two features: $\pm N$ and $\pm V$. Prepositions correspond to $[-N, -V]$. Though this typically leads authors to include prepositions on the list of major lexical categories, it is done with some ambivalence. For instance, Chomsky (1981) presents this feature system, but then says that only the nouns, verbs, and adjectives are lexical categories. But this appears to be an off-hand remark, having more to do with traditional terminology than a really substantial claim. Chomsky (1986a) presents the same feature system, and counts all four categories, including prepositions, as lexical. In work taking the distinction between lexical and functional categories more seriously, we see Baker (2003) arguing explicitly that prepositions are functional, and Hale and Keyser (2002) highlighting their mixed status. The reason prepositions are such a difficult case is that some prepositions, or some occurrences of some prepositions, seem to play an entirely grammatical role, while some do not. The prepositions expressing thematic properties, such as dative *to* expressing recipienthood, seem functional in nature. Other prepositions, including prepositions expressing spatial position like *in*, may appear to be lexical.

When it comes to metaphor, we get the same sorts of mixed results. Thematic prepositions like dative *to* do not generate any metaphorical interpretations. With prepositions like *in*, the situation is less clear. There are abstract uses of *in* like:

- (i) John is in trouble.

These are often glossed as having metaphorical interpretations of the spatial relation of being in (cf. Quirk *et al.*, 1985). However, it is not entirely obvious whether this is metaphor, or rather a highly abstract meaning of *in*. (*To* itself has locative uses implying directionality, for which similar questions can be raised.) At least, in these cases, we get mixed results about metaphor just where we have a mixture of functional and lexical characteristics.

⁸ There are a number of more technical ways that lexical and functional categories may be distinguished. Semantically, functional categories tend to have the semantics of operators. Syntactically, functional categories do not discharge or assign theta-roles, while they do bear agreement features (cf. Fukui, 2001). In some frameworks, not relating to theta-roles and having the semantics of operators come to pretty much the same things. The metaphor probe can work in conjunction with these more theoretical ideas, as it gives us a more intuitive test, and a more intuitive indication of what is different about the semantics of functional categories. Even so, the speculative suggestion I shall make in section 5 may indicate that theta-roles and metaphor have a great deal to do with each-other.

⁹ For a survey of philosophical issues about metaphor, see Reimer and Camp (2006).

¹⁰ See Camp (2006) for discussion and references.

¹¹ *Concept* may not be the best term to use here, for two reasons. First, not every view of metaphor cognition sees the units of cognition as playing the same roles that concepts traditionally play. Second, many views of metaphor work with larger-scale elements of cognition (such as scripts or schemas). Even so, these differences will not be important for the little I shall say about psychological models of metaphor cognition, so talking about

concepts will be harmless.

¹² My debt to Camp (2006) in this section should be very obvious.

¹³ For a taste of the debate about semantics in the language faculty, see Larson and Segal (1995) and Pietroski (2003).

¹⁴ I have followed the proposal of Moro (1997). It might be surprising to some how much controversy there is over the syntax of such an apparently simple sentence. For an overview, in a more recent syntactic framework, see Mikkelsen (2005).

¹⁵ For the moment, we may avoid the issue of how lexical items link to concepts. I shall discuss some aspects of the lexicon in section 4, but I shall generally skirt the question of whether lexical entries themselves contain conceptual information, or rather contain some kind of pointers to concepts. The interface model we are considering can work with either.

¹⁶ A tempting speculation that would fit with the simple model is that though the meanings of expressions like *more* and *bigger* are related, the genuine lexical categories include links to the broader conceptual system, whereas the functional ones have home-grown contents provided by the language faculty itself. However, the reasons the simple model fails I shall present in section 4 show that even if this speculation is correct, it cannot account for the interface behavior we see. At least some home-grown linguistic contents do get passed to the metaphor system.

¹⁷ Though most anything in lexical semantics is controversial, this picture has emerged from a wide range of research. Here is a representative statement, “The argument I will make is that semantic properties of predicates divide into two fundamentally different kinds of information. I suggest that the division corresponds to the distinction between information that is linguistically analyzed and information that, while it may be cognitively analyzed, is linguistically atomic. The argument is based on the idea that some meaning components have grammatical life, and some are linguistically inert. There is a sense in which this position is a distillation of a general research trend . . .” (Grimshaw, 2005b).

¹⁸ For more critical discussion of the theory of aspectual classes, see Levin and Rappaport Hovav (2005) and Verkuyl (1993).

¹⁹ Sometimes events and states are distinguished. If so, then verbs describe a broader category of ‘eventualities’ including both.

²⁰ I have put this in terms of how a verb depicts an event, and so put the difference in the verb’s meaning rather than the event itself. Whether or not there is a difference between events answering to properties like telicity depends on just how finely events are individuated. Fine-grained event individuation will see the differences we are discussing here as differences between events, while coarse-grained event individuation will not. As our main interest here is in the meanings of verbs, we need not take a stand on this issue.

²¹ Many theories of aspectual classes recognize another feature of punctuality. With it, they identify an additional class of *semelfactives* like *kick* and *wink*, which are near-instantaneous, but not telic. See Smith (1997).

²² See Rothstein (2004), who builds on work of Krifka (e.g. Krifka, 1998) and Landman (1992).

²³ Ultimately, Rothstein (2004) proposes a more refined analysis of accomplishments.

²⁴ Some approaches to lexical meaning make structural elements even more syntactic. See, for instance, Hale and Keyser (2002).

²⁵ In a few tests I ran, I found that about half of my informants were able to interpret

extended metaphors even when the temporal information in the surrounding text and the aspectual structure of a target sentence conflict. Many who could not find the metaphors infelicitous.

²⁶ For more on the internal structure of progressive achievements, see Rothstein (2004).
²⁷ The implicature is not generally present in progressive achievements. We do not see it in:

- (i) a. John was dying when I looked at him.
 b. Mary stava arrossendo quando la guardai.
 Mary was blushing when I looked at her.

However, this does not have a punctually locating reading. Rather, it locates a stage of the accomplishment, i.e. a stage of a *DO*ing rather than a *BECOMING*. This does not trigger the implicature. The analysis of progressive achievements of Rothstein (2004) does include a *BECOME* operator in a much more complex frame for them, but it relates to the culmination of the accomplishment, which presumably does not support the inference that the speaker is responsible for the becoming.

²⁸ Example from http://science.nasa.gov/headlines/y2004/07jul_bluemoon.htm.

²⁹ Thanks to Adam Sennet for suggesting this line of argument. Stern (2000) argues that the badness of sentences like (32), which employ conjunction reduction to mix metaphorical and literal interpretations of a single expression, is evidence that metaphor should be treated semantically. Camp (2005) argues against this interpretation of the data; I take no position on the issue here.

³⁰ One alternative interpretation of (33) is that it literally states a fact about the patient's skin color, and thereby implicates something about the cause of that state. Even so, it does sound on par with other metaphors to my ear, and I suspect that other examples are even less vulnerable to this worry.

³¹ In related work, Asher and Lascarides (2001) argue that metaphor processing itself works on highly structured lexical entries, represented in their framework as typed feature structures. They argue that there is a substantial constraint on metaphorical interpretations derived from these structures, which shows that certain structural features of verbs cannot be changed in metaphorical interpretation. In particular, they argue that the structural features that distinguish the class of change of location verbs are preserved in metaphor. Likewise Stern (2000) notes that the thematic structure of a verb is preserved in metaphorical interpretation, and that thematic structure can trigger metaphorical interpretation of a verb's arguments.

³² For discussion of nominals and adjectives, see Grimshaw (1990) and Larson (1998).

³³ This might also be an argument for Higginbotham's thesis about event arguments. Those of us coming from a more Montagovian tradition are rather cautious about this claim, but at least indirectly, the evidence I have been discussing here may well support it. It also would bear out the idea, mentioned briefly in section 2, that assignment and discharge of theta-roles is the key distinguishing feature between lexical and functional categories.

References

- Asher, N. and A. Lascarides, 2001. Metaphor in discourse. In P. Bouillon and F. Busa (eds.), *The Language of Word Meaning*, pp. 262–289. Cambridge: Cambridge University Press.
- Baker, M. C., 2003. *Lexical Categories*. Cambridge: Cambridge University Press.
- Black, M., 1962. Metaphor. In *Models and Metaphors*, pp. 25–47. Ithaca, N.Y.: Cornell University Press.
- Camp, E., 2005. Critical study of Josef Stern's *Metaphor in Context*. *Noûs* 39:715–731.
- , 2006. Metaphor in the mind: The cognition of metaphor. *Philosophy Compass* 1:154–170.
- Carston, R., 2002. *Thoughts and Utterances*. Oxford: Blackwell.
- Chomsky, N., 1970. Remarks on nominalization. In R. A. Jacobs and P. S. Rosenbaum (eds.), *Readings in English Transformational Grammar*, pp. 184–221. Waltham: Ginn and Company.
- , 1980. *Rules and Representations*. New York: Columbia University Press.
- , 1981. *Lectures on Government and Binding*. Dordrecht: Foris.
- , 1986a. *Barriers*. Cambridge: MIT Press.
- , 1986b. *Knowledge of Language*. New York: Praeger.
- Davidson, D., 1978. What metaphors mean. *Critical Inquiry* 5:31–47. Reprinted in Davidson (1984).
- , 1984. *Inquiries into Truth and Interpretation*. Oxford: Oxford University Press.
- Dowty, D. R., 1979. *Word Meaning and Montague Grammar*. Dordrecht: Reidel.
- Fodor, J. A., 1983. *The Modularity of Mind*. Cambridge: MIT Press.
- Fukui, N., 2001. Phrase structure. In M. Baltin and C. Collins (eds.), *Handbook of Contemporary Syntactic Theory*, pp. 374–406. Oxford: Blackwell.
- Gentner, D. and P. Wolff, 1997. Alignment in the processing of metaphor. *Journal of Memory and Language* 37:331–355.
- Glucksberg, S. and B. Keysar, 1993. How metaphors work. In A. Ortony (ed.), *Metaphor and Thought*, pp. 401–424. Cambridge: Cambridge University Press, 2nd edn.
- Grimshaw, J., 1990. *Argument Structure*. Cambridge: MIT Press.
- , 2005a. Extended projection. In *Words and Structure*, pp. 1–73. Stanford: CSLI Publications.
- , 2005b. Semantic structure and semantic content in lexical representation. In *Words and Structure*, pp. 75–89. Stanford: CSLI Publications.
- Hale, K. and S. J. Keyser, 2002. *Prolegomenon to a Theory of Argument Structure*. Cambridge: MIT Press.
- Higginbotham, J., 1985. On semantics. *Linguistic Inquiry* 16:547–593.
- , 1987. The autonomy of syntax and semantics. In J. L. Garfield (ed.), *Modularity in Knowledge Representation and Natural-Language Understanding*, pp. 119–131. Cambridge: MIT Press.
- , 2000. On events in linguistic semantics. In J. Higginbotham, F. Pianesi, and A. C. Varzi (eds.), *Speaking of Events*, pp. 49–79. Oxford: Oxford University Press.
- Krifka, M., 1998. The origins of telicity. In S. Rothstein (ed.), *Events and Grammar*, pp. 197–235. Dordrecht: Kluwer.
- Lakoff, G., 1993. The contemporary theory of metaphor. In A. Ortony (ed.), *Metaphor and Thought*, pp. 202–251. Cambridge: Cambridge University Press, 2nd edn.
- Lakoff, G. and M. Johnson, 1980. *Metaphors We Live By*. Chicago: University of Chicago

- Press.
- Landman, F., 1992. The progressive. *Natural Language Semantics* 1:1–32.
- Larson, R. K., 1998. Events and modification in nominals. *Proceedings of Semantics and Linguistic Theory* 8:145–168.
- Larson, R. K. and G. Segal, 1995. *Knowledge of Meaning*. Cambridge: MIT Press.
- Levin, B. and M. Rappaport Hovav, 1995. *Unaccusativity: At the Syntax-Lexical Semantics Interface*. Cambridge: MIT Press.
- , 2005. *Argument Realization*. Cambridge: Cambridge University Press.
- McClure, W., 1990. A lexical semantic explanation for unaccusative mismatches. In K. Dziwirek, P. Farrell, and E. Mejías-Bikandi (eds.), *Grammatical Relations*, pp. 305–318. Stanford: CSLI Publications.
- Mikkelsen, L., 2005. *Copular Clauses*. Amsterdam: John Benjamins.
- Moro, A., 1997. *The Raising of Predicates*. Cambridge: Cambridge University Press.
- Ortony, A., 1979. Beyond literal similarity. *Psychological Review* 86:161–180.
- Pietroski, P., 2003. The character of natural language semantics. In A. Barber (ed.), *Epistemology of Language*, pp. 217–256. Oxford: Oxford University Press.
- Quirk, R., S. Greenbaum, G. Leech, and J. Svartvik, 1985. *A Comprehensive Grammar of the English Language*. Harlow: Longman.
- Recanati, F., 2004. *Literal Meaning*. Cambridge: Cambridge University Press.
- Reimer, M. and E. Camp, 2006. Metaphor. In E. Lepore and B. C. Smith (eds.), *Oxford Handbook of Philosophy of Language*, pp. 845–863. Oxford: Oxford University Press.
- Rothstein, S., 2004. *Structuring Events*. Oxford: Blackwell.
- Ryle, G., 1949. *The Concept of Mind*. London: Hutchinson.
- Searle, J. R., 1979. Metaphor. In *Expression and Meaning*, pp. 76–116. Cambridge: Cambridge University Press.
- Smith, C. S., 1997. *The Parameter of Aspect*. Dordrecht: Kluwer, 2nd edn.
- Sperber, D. and D. Wilson, 1998. *Relevance*. Cambridge: Harvard University Press.
- Stern, J., 2000. *Metaphor in Context*. Cambridge: MIT Press.
- Van Valin, R. D., Jr., 2005. *Exploring the Syntax-Semantics Interface*. Cambridge: Cambridge University Press.
- Vendler, Z., 1967. *Linguistics in Philosophy*. Ithaca: Cornell University Press.
- Verkuyl, H. J., 1993. *A Theory of Aspectuality*. Cambridge: Cambridge University Press.
- Walton, K. L., 1993. Metaphor and prop oriented make-believe. *European Journal of Philosophy* 1:39–56.