TABLE OF CONTENTS

LIST OF FIGURES

INTRODUCTION: WHY DO WE NEED A THEORY OF METAPHOR?

As I think of it, and as many other people think of it, metaphor is fundamentally a linguistic phenomenon. There are other uses of the term 'metaphor' — for example, people sometimes talk about visual or musical metaphors. And as we'll see, metaphor is intimately bound up with certain distinctive, complex patterns and processes of thought. However, in this dissertation I will be concerned with metaphor as feature of language exclusively. If we can get clear on how the linguistic case works, then we can turn to these other cases of 'metaphor', and see just what their relation is to the linguistic phenomenon. Most likely, they will turn out to depend on much the same complex patterns and processes of thought as linguistic metaphors do.

Within the realm of language, we use metaphor in a wide variety of ways. We use it to talk about the world in both familiar and innovative terms: some metaphorical ways of speaking are so familiar that we barely notice that they are metaphorical; other metaphors are so wild and novel that we puzzle for a long time over what they could mean. And we use both familiar and innovative metaphors in a wide range of contexts, ranging from everyday conversation to literature and scientific theorizing. Such wide variety makes it especially challenging, but also vitally important, to determine what all these ways of speaking have in common, and why we should be so drawn to speaking in this way.

Unfortunately, philosophers have not risen well to this challenge. As Simon Blackburn says, "nobody would claim that the study of metaphor has been one of analytic philosophy's brighter achievements."[1] Philosophers tend to treat metaphor as the black sheep of language — the "poor relation of proper judgment."[2] It is marginalized for its unruly behavior, and acknowledged only grudgingly as a member of the family. When mainstream analytic philosophers do treat metaphor at all, they tend to think about it just long enough to come up with a general theory, but not long enough to delve into the details of how metaphor actually works. The result has been a series of mutually conflicting and individually less-than-convincing papers by some of philosophy's best minds. Other theorists have taken

---

[1] *Spreading the Word*, p. 180.
[2] *Spreading the Word*, p. 171.

metaphor more seriously, but they have tended to become so transfixed by the wonders of metaphor — its expressive powers, its pervasiveness, its profundity — that they neglect the insights of more mainstream philosophy, and turn into evangelists for their cause.

Metaphor causes all this trouble because it straddles so many of the usual boundaries: between language and thought, between semantics and pragmatics, between rational thought and mere causal association.   In particular, it challenges the usual philosophical picture of communication as the maximally efficient transmission of information. Metaphor seems to involve speaking circuitously where direct expression is often possible. And metaphor is often used to produce non-informational, imaginative and emotional responses. But on the usual picture of the mind as a rational representational faculty, and of communication as a rational cooperative enterprise, the fact that we use metaphor so often to accomplish such strange ends is difficult to explain.

Nonetheless, a complete philosophical account of the way language works must include a theory of metaphor. Metaphor is an important communicative tool in its own right. In addition, our inability to make sense of such a pervasive and natural feature of the way we talk suggests that the traditional picture may be distorted in other ways as well. In fact, I will argue that uncovering the resources necessary for an adequate theory of metaphor will open our eyes to a range of other linguistic and cognitive phenomena that we also tend to neglect, and thereby reveal other avenues for philosophical investigation.

In chapter 1, I survey what I take to be the three main types of a theory of metaphor. Philosophers of language divide the terrain of linguistic meaning into three major areas: semantics, or phenomena related to sentence meaning; pragmatics, or phenomena related to speaker meaning; and perlocutionary effects, or the other, further effects which a speaker's utterance brings about in a hearer. The three main types of theory of metaphor each claim that metaphor belongs to one of these linguistic domains.

On my view, metaphor is a pragmatic phenomenon: in speaking metaphorically, speakers exploit the conventional meanings of their words to undertake speech acts with distinct propositional contents. Whether an utterance is metaphorical at all, and if so, what its metaphorical content is, depend more on

the speaker's intentions than on the conventional meaning of the words uttered. Because I think the speaker's intentions play such a crucial role here, I also think that an account of metaphor must uncover the psychological principles on which metaphor operates: what patterns of thought enable hearers to recover the speaker's intended content? What features of these patterns distinguish metaphor from other uses of language?

Theorists who treat metaphor as a semantic phenomenon often do so because they have identified important, distinctive ways in which metaphor differs from typical cases of implicature, which is the paradigmatic pragmatic phenomenon. Metaphorical meaning is more intimately connected to conventional word meaning than implicature is, and metaphorical utterances behave differently in conversation than implicatures often do. These phenomena call for explanation, and a simple theory on which metaphor is treated simply as a brand of implicature will not do. I argue, though, that the evidence marshaled by these theorists can in fact be incorporated into a more nuanced and sophisticated pragmatic account.

Theorists who think metaphor belongs to the third domain focus on the rich cognitive and imaginative effects that novel, poetic metaphors bring about. Metaphors make us *see* one thing *as* another, they say, and thereby make us think about that thing in a different 'light'. Because these effects are largely non-propositional, they aren't the sort of thing that a speaker might *mean*. These theorists therefore conclude that metaphor cannot be a pragmatic phenomenon. I agree that 'seeing-as' or 'seeing under an aspect' is an essential component of metaphor. But where these theorists take 'seeing-as' to be the sole, final *product* of metaphorical comprehension, I propose that it is the *means* by which speakers intend that their hearers recover their metaphorical, propositional content. The rest of the dissertation develops this proposal into a theory of metaphor.

Any theory that invokes 'seeing-as' must be explicit about just what this notion amounts to. This is especially imperative given that talk of 'seeing-as' is itself metaphorical: Juliet is not visually presented to us when we 'see' her as the sun; and we cannot see life, even in our mind's eye, as "but a walking shadow." In chapter 2, I examine the perceptual phenomenon of seeing-as in order to elucidate its

counterpart in thought. Crucially, when we see something in a new way, a concept or thought re-structures, and does not merely exist alongside, the perception itself. To make sense of the analogous phenomenon in thought, we thus need to identify an appropriate type of structured mental representation, and to determine how it too can be restructured. In chapter 3, I develop the notion of a *characterization* of an individual or kind. Characterizations are distinguished from concepts in at least three key respects: they include additional experiential and encyclopedic information; they include non-truth-conditional elements, such as emotions; and their elements vary in their prominence and centrality with respect to one another. In chapter 4, I describe how an aspect *F*, can re-structure one's characterization of an individual *a*, by highlighting features in *a*'s characterization that can be matched in certain ways to prominent or central features in *F's* characterization.

In chapter 5, I put this account of aspectual thought to work on a theory of metaphor. In particular, I show that my account can treat the full range of metaphors, from 'ordinary' to 'poetic'. This single distinction is in fact the product of at least four independent variables: a metaphor's conversational weight, the ease with which its content can be communicated by other linguistic means, the richness and novelty of the aspect it generates, and the speaker's commitment to the aptness of the generating aspect itself. Attending to these variables allows us to understand why ordinary conversational metaphors are efficient and sometimes essential vehicles for communicating determinate contents. It allows us to understand how rich, poetic metaphors can be invitations to an open-ended contemplation of the subject's overall character. And it allows us to do both of these things without abandoning a unified theory of metaphor.

In chapter 6, I step back from my own account of metaphor to address two related questions. First, I deploy the resources and insights developed in the course of the previous chapters to address the vexed question of whether metaphors can be paraphrased. So little progress has been made on this question in part because there has been no real agreement about what an adequate paraphrase should accomplish, and in part because theorists have taken different classes of metaphor to be paradigmatic. I conclude that an adequate paraphrase can be articulated for most metaphors, albeit with quite a lot of

effort. A paraphrase can and should state all the commitments the speaker has undertaken about the subject under discussion. A paraphrase must, however, leave out what is most distinctive about metaphor: the cultivation of the aspect itself, as the means to determining the speaker's claim. And without the aspect, much of the cognitive power and interest of the original utterance is inevitably lost.

The second question I ask concerns what the investigation into metaphor can teach us about philosophy of mind and of language more generally. I conclude that it forces us to attend to patterns of thought which have been largely ignored because they fail to fit the traditional model of the mind as a *calculus ratiocinatur*, but which nonetheless deserve philosophical consideration. Once we recognize the role these patterns of thought play in the comprehension of metaphor, we can also see that they often play a role in literal communication as well, although that role is usually less palpable and extensive. What is essentially distinctive about metaphor is that in it, aspectual thought serves as the means to determining the content of the speaker's primary communicative intention. But in other respects, the gap between literal and metaphorical communication is narrower than we might otherwise have thought.

x

CHAPTER 1: TOWARD A PRAGMATIC THEORY OF METAPHOR

I think we should expect a theory of metaphor to accomplish at least three tasks. First, it should locate metaphor within the broad range of linguistic phenomena. Second, it should describe the typical effects of metaphor, with special attention to the effects that distinguish it from other, similar phenomena: *what* does metaphor do, and what is special about it? How does it differ from literal speech, or from, say, insinuation, irony, or metonymy? Third, a theory of metaphor should explain, at least in outline, how those effects are achieved: *how* does metaphor manage to do what it does? In this chapter, I set the stage for my own theory of metaphor by outlining the main candidate answers to the first question — of where to locate metaphor — and then seeing how theories which answer this question differently in turn offer different answers to the latter two questions. This survey will then allow me to sketch the form of my own account in a quick and efficient way.

Within the realm of linguistic phenomena, philosophers typically distinguish find three broad regions to which metaphor might belong: semantics, pragmatics, and the further uses to which language is put — the perlocutionary realm, in Austin's terminology. Unfortunately, none of these terms has a fully settled definition. In some cases the debate about how to define them is merely a war of words; in others it is symptomatic of important differences in overall theoretical commitment and outlook. I will use them in the following way, which I think is both moderate and relatively common. Readers who disagree with my way of using these terms should be able to translate my discussion into their preferred taxonomy.

Semantic phenomena are features that an utterance exhibits by virtue of the *conventional meanings* of its constituent words and their mode of combination. This includes features that depend upon the particular context of utterance, so long as the conventional meaning of the sentence uttered determines that those features are relevant. Thus, the following are all semantic facts, on this definition: the fact that utterances of 'I' in English conventionally refer to the utterance's speaker, the fact that *this* utterance of 'I' conventionally refers to me, and the fact that this utterance of

(1) I am running

1

conventionally expresses the proposition that this person, Liz, is running at this time, t_0. On my use of the term 'semantic', semantic facts concern what the *sentence token* means. Notice that on this definition, semantic facts can and typically do include more than just a specification of the truth-conditional content expressed by an utterance. They also include at least general conventional constraints on the utterance's illocutionary force and any conventional non-truth-conditional expressive content. So, the fact that (1) has the illocutionary force of an assertion, and the fact that

(2) That damn Liz is out running

expresses a derogatory attitude toward Liz, also count as semantic facts, on this definition. I will employ talk of 'what a speaker *says*' to include all and only what the linguistic conventions fix as the meaning of her utterance. Thus, in my terminology here, 'what a speaker says' just is what her utterance semantically expresses.[1]

Pragmatic phenomena are features that an utterance exhibits by virtue of the speaker's intention to get the hearer to recognize what she's trying to communicate to him. The speaker's pragmatic intentions are satisfied just in case the hearer recognizes what those very intentions are. Thus, the fact that a particular utterance of

(3) You're late

is a reprimand for the hearer's tardiness is a pragmatic fact; so is the fact that a particular utterance of

(4) It's getting late

is a request to the hearer to leave soon. Pragmatic facts concern what the *speaker means* in uttering the sentence she does. Often, even when an utterance is literal and direct, the speaker's meaning will be richer and more determinate than the meaning of the sentence uttered. In particular, both the illocutionary force and the intended truth-conditions of even literal utterances are typically more specific than what is determined by the conventional sentence meaning. So, for instance, a speaker may utter

(5) Open the door

---

[1] This is the strict, technical sense of 'says', which can come apart from the loose use of 'says' in ordinary speech. I return to this distinction in §6.1.

intending it not just as some sort of directive, but specifically as a demand, and intending to demand not just that the hearer open the door in some way or other, but that he open it by drawing the lock and turning the handle.[2]   Because there is typically such a pervasive difference between what a speaker says and what a speaker means even for literal direct discourse, pragmatic meaning plays an essential role in all linguistic comprehension. Indeed, even determining that the speaker intends her utterance literally and directly requires a hearer to engage in pragmatic processing.[3]

Perlocutionary phenomena concern the further effects which the speaker achieves by making her utterance, consequent on the hearer's recognition of her communicative intentions.[4] So, for instance, the fact that the utterance of (3) above makes the hearer frightened, that it makes him show up on time in the future, or that it makes the speaker seem more authoritative, are perlocutionary facts.

Given this rough schema, we should expect three broad types of theories of metaphor. A theory might classify metaphor as a semantic, a pragmatic, or a perlocutionary phenomenon. There are, of course, important differences among theories within these categories, but this broad classification provides us with an overview of the territory. I'll consider each type in turn, focusing on just one or two representative versions of each type.

1.1: *Semantic Theories*

1.1.1: *'Old-School' Semanticists*

'Old-school' semanticists draw on the intuition that metaphorical meaning is intimately bound up with the conventional meanings of the words uttered, in a way that is not true of typical pragmatic phenomena. Utterances of different sentences which express largely the same propositional content using different words will almost always produce different metaphorical effects. Consider here the difference in metaphorical effect between utterances of

(6) After his oral exam, George stood on stage like a steed breathing hard after a quick gallop.

---

[2] Cf. John Searle, "Literal Meaning."
[3] Cf. Kent Bach and Robert Harnish, *Linguistic Communication and Speech Acts*, p. 63.
[4] Cf. J. L. Austin, *How to Do Things with Words*, p. 121.

and

(6') After his oral exam, George stood on stage like a nag breathing hard after a quick gallop.

(6) suggests that George did well in his exam, and perhaps even enjoyed it, while (6') suggests just the opposite. **Conversational implicatures**, by contrast, typically depend more on the propositional content itself than on how that content is expressed. (This is what Grice calls the "nondetachability" of conversational implicatures, because the implicature cannot be detached from the content.)[5] The fact that the particular words uttered play such a crucial role in metaphor motivates semantic theories on which the words do *all* the work of fixing metaphorical meaning. What is distinctive about 'old-school' semanticists in particular is that they try to build all the resources for delivering the metaphorical meaning into the semantics itself. By contrast, as we'll see, the 'new semanticists' allow context to play an important role as well.

More specifically, old-school semanticists maintained that a metaphorical sentence is *semantically deviant*, in that it is logically absurd or self-contradictory on a literal reading. This deviance produces a 'metaphorical twist,' in which the meanings of at least some of the sentence's constituent words are altered. How they are altered, and so what the metaphorical meaning of the complete sentence turns out to be, is itself determined semantically. Where old-school semanticists differed significantly among themselves was in *how* they thought the semantics determined this new meaning. We can identify two important strands within the old-school semantic tradition.

The first strand attempts to derive a sentence's metaphorical meaning directly from the definitions of its constituent terms; proponents of this sort of view include Uriel Weinreich,[6] Samuel Levin,[7] and L. Jonathan Cohen and Avashai Margalit.[8] These theorists come from the tradition of lexical semantics, on which the meaning of a term is given by its lexical entry, consisting of selection restrictions

---

[5] Typically, though by no means always — 'manner' implicatures are a general exception here. Cf. Paul Grice, "Logic and Conversation," p. 39. Nonetheless, I think something like this intuition has driven many theorists who take a semantic view of metaphor.
[6] "Explorations in Semantic Theory."
[7] *The Semantics of Metaphor*.

and semantic markers.[9] (Selection restrictions specify the syntactic and semantic conditions that must be met by a larger linguistic context if the term in question can be inserted into that context. Semantic markers specify the attributes that the term contributes to the whole sentence's truth-conditions.) The grammar then provides projection rules for combining the meanings of the constituent terms into successively larger phrases, culminating in a reading for the whole sentence.

The assumption is that no such reading will be delivered for metaphorical sentences, because the constituent terms will contain incompatible selection restrictions, and the normal process of composition will therefore be blocked. So, for instance, the sentence

(7) The stone died

 is taken to be literally nonsensical, because a semantic marker for 'stone' is <-animate>, while a selection restriction for 'to die' is <+animate>. ==Metaphorical interpretation then consists of either transferring or deleting markers and restrictions so that composition can proceed.== Metaphorical meanings will typically be highly multiply ambiguous. This is because we can 'metaphorize' multiple terms in the sentence, because each of those terms will themselves possess multiple markers and restrictions, and because each of those markers can either be transferred to another term or else deleted altogether. Each different modification will produce a different metaphorical reading. Thus, for (7), we could delete the <-animate> restriction on 'the stone', thereby making the term's denotation a stone-like living thing. Or alternatively, we could delete <+animate> from 'to die', to specify that the stone ceased to exist.[10]

---

[8] "The Role of Inductive Reasoning in the Interpretation of Metaphor."

[9] See Katz and Fodor, "The Structure of a Semantic Theory," for the classic statement of this theory of semantics.

[10] Samuel Levin offers six such 'construal rules', three on which we modify the meaning of the noun phrase, and three the verb phrase. We may *transfer* a semantic marker from the verb to the noun (or vice versa), and then either (1) *disjoin* or (2) *conjoin* this marker with others already in the noun's lexical entry. We may also (3) *replace* a semantic marker in the noun with one from the verb (or vice versa). Using the sentence 'The stone died', Levin claims that we get the following six construals:
   Noun modified:
     (1) The concreted physical object that's either mineral or human died.
     (2) The concreted personified stone died.
     (3) The concreted human (that is, unfeeling or stupid person) died.
   Verb modified:
     (4) The stone ceased to exist.
     (5) empty interpretation
     (6) The stone disintegrated.

An account of metaphor of this sort would be quite compelling, if it worked. We would at one fell swoop assimilate metaphor within semantics by utilizing only resources that are already required by the grammar (the lexical entries); we would cleanly differentiate metaphor from literal meaning and isolate a necessary and sufficient criterion for metaphoricity (semantic deviance); and we would provide a relatively clear and detailed rule for interpretation which captured the initial intuition that metaphorical meaning depends in an essential way on literal meaning (the construal rules). However, matters do not turn out to be so neat and simple.

Námitky proti sémantickým teoriím:

First, not all metaphorical sentences are semantically deviant, let alone logically absurd or nonsensical. Ted Cohen[11], Michael Reddy[12] and others have delighted in pointing out that sentences like

(8) No man is an island

are semantically unimpeachable but can still be used metaphorically; that sentences like

(9) The rock is becoming brittle with age

have perfectly normal contexts of literal application, but can also be used metaphorically in other contexts (such as to describe a dottering professor emeritus); and that sentences like

(10) Moscow is a cold city

can be literally *and* metaphorically true, and even conversationally appropriate, in one and the same conversational context.[13]

So semantic deviance cannot be a necessary condition for metaphoricity. A significant part of the semantic account's clean explanatory power is therefore undermined. We will at least need to supplement that account with a theory of what makes an utterance metaphorical in the first place, since our original condition has turned out to be unworkable. Given the examples just cited, it would seem intuitively likely

---

[11] "Figurative Speech and Figurative Acts."
[12] "A Semantic Approach to Metaphor."
[13] Patti Nogales (*Metaphorically Speaking*, p. 2) offers this nice example to make a similar point:
   Dena: The man lives without a doorbell or any windows.
   Diane: He's hard to get in touch with, huh?
   Dena: No, seriously, his house doesn't have a doorbell or any windows.
Here both the literal and metaphorical construals are *prima facie* contextually appropriate; the hearer chooses the one that attributes the more plausible proposition, but the speaker actually intends to claim that the more unusual state of affairs obtains.

==that whether a given utterance is to be interpreted literally or metaphorically is a pragmatic matter: a matter of how the speaker intends to use her words in the particular context of utterance.== However, perhaps we could still retain the semantic account as a theory of metaphorical *meaning*. Perhaps, that is, the semantics could determine the range of a sentence's possible readings, both literal and metaphorical. All of these might then be handed over to the pragmatic processing module, which would then select one from among the semantically determined options as the speaker's intended meaning.[14]

==It is precisely in the determination of possible meanings, however, that the lexical theory faces its greatest difficulty.== The fundamental assumption of lexical theories of metaphor is that the resources for determining metaphorical readings are drawn entirely from the lexical entries for the sentence's constituent terms. But these resources are often insufficient for generating the intuitively appropriate metaphorical interpretations of even relatively straightforward, prosaic metaphors. To remedy this insufficiency, lexicalists end up incorporating features we would not otherwise think of as semantic within their postulated lexical entries. Thus, Levin is led to maintain that "rapacity, ., etc" are part of the lexical entry for 'wolf', in order to generate the appropriate metaphorical interpretation of

   (11) Man is a wolf.[15]

Because Levin simply adopts a standard Katz-Fodor style selection-restriction account, he provides no criterion of his own for when a feature should be included within a lexical entry. (Indeed, he claims that Katz has no clear criterion of this sort, either.[16]) We thus have no principled method for deciding whether the entry for 'wolf' should include these features on Levin's own terms. But we can certainly say that these are not the sort of features we would intuitively treat as definitional of being a wolf, or even as necessarily applicable to wolves, as semantic theories of this traditional sort usually require.

---

[14] In this vein, Cohen and Margalit abjure the requirement of semantic deviance, and maintain instead that "we distinguish instead between sentences that admit only metaphorical interpretations, like *That old man is a baby*, and sentences that admit both metaphorical and non-metaphorical ones, like *Tom is a baby*." ("The Role of Inductive Inference," p. 738).
[15] *The Semantics of Metaphor*, p. 54.
[16] *The Semantics of Metaphor*, pp. 95-99.

Cohen and Margalit, by contrast, do provide an explicit criterion for a feature's having a semantic status. In rough terms, they include within the definition of a term F all and only those features Φ_1, $\Phi_2,...\Phi_n$ the possession or non-possession of which by x could falsify the hypothesis that x is F.[17] This criterion is intended to deliver relatively rich and inclusive definitions — certainly more inclusive than those of the traditional Katz-Fodor theory. But even here, Cohen and Margalit are led to include as semantic some features that fail to meet their stated criterion. Thus, they end up treating small size and mental incapacity as semantic features of 'baby', in order to deliver the appropriate metaphorical reading of 'baby airplane'. They likewise include being underprivileged as a semantic feature of 'negro' in order to deliver the intuitive reading of

(12) The poor are the negroes of Europe.[18]

Perhaps the explicit criterion for semantic status could be relaxed even further, so that these and other features intuitively relevant to interpreting metaphor *would* count as semantic after all. This would seem to be a rather bad idea, especially given the examples Cohen and Margalit provide. But if we do relax the criterion in this way, then we will face a new problem. We will have included so many features within our lexical entries that linguistic interpretation — both literal *and* metaphorical — will effectively turn out to be a pragmatic matter after all. Semantic ambiguity climbs exponentially with the addition of each lexical feature. The number of readings delivered by a suitably enriched semantics will thus be astronomically large.[19] Although the resolution of ambiguity is often considered not to be a distinctively pragmatic matter, still the number of available readings will end up being so overwhelmingly great that lexical input will be only weakly relevant to the output of actual linguistic comprehension. The semantic theory would not be inconsistent or demonstrably false, but it would be practically useless.

---

[17] "The Role of Inductive Reasoning," p. 728.

[18] "The Role of Inductive Reasoning," p. 735.

[19] So, for instance, on Levin's theory, assuming an average of twenty total selection restrictions and semantic markers (a conservative estimate given the range of features required just to accommodate the particular examples cited), the simplest subject-predicate sentence, like 'The stone died' or 'Man is a wolf' would have a minimum of 120 readings. Even this accounts only for meaning changes that involve just a single semantic marker; often, though, the intuitively appropriate reading will require deleting or transferring multiple markers, as with Cohen and Margalit's 'baby airplane' example. This will drive the number of readings even higher.

This difficulty is fundamental for old-school semantic theories of the selection-restriction type. **The second** strand of old-school theorists, the 'connotationists', try to avoid it by distinguishing between those features that are relevant to fixing a term's denotation and those that are relevant to its connotation. Metaphorical meaning, they claim, is primarily a function of the latter. Monroe Beardsley[20] (perhaps along with I. A. Richards[21] and Max Black[22]) is the most prominent defender of this sort of view; I will focus on Beardsley's account. 'Connotationists' tend to be much more concerned than lexicalists about explaining the rich, nuanced effects of poetic metaphors.

The leading idea once again, is that there is a "logical absurdity"[23] embedded in the sentence at the literal level, which prevents the predicate or 'modifier' (that is, the metaphorical term) from combining with the subject.[24] This "inherent tension" again forces a "twist of meaning"[25] in the predicate: as Beardsley says, "the price it [the predicate] pays for admission to this context is that it function there to signify only its connoted characteristics."[26] If those connoted characteristics include features that could plausibly, though perhaps not truly, be attributed to the subject, then the sentence is metaphorical; otherwise, it is nonsensical.[27] Beardsley's position seems to be that the predicate attributes *all* of those plausible connoted characteristics to the subject, and so that a metaphorical sentence itself is

---

[20] *Aesthetics*, "The Metaphorical Twist."

[21] *The Philosophy of Rhetoric.* Empson (*The Structure of Complex Words*) expands and clarifies Richards' view.

[22] "Metaphor," "How Metaphors Work: A Reply to Donald Davidson." Because Black does not argue explicitly that metaphor is a distinctively semantic phenomenon, I will not take up his view directly here. Several remarks suggest, however, that he does intend his theory to be a semantic one. Both Richards and Black are much more interested in how the metaphorical term's connotations 'interact' with the subject term to produce a new meaning than in whether that meaning should be classified as semantic or pragmatic. I discuss Black's view, without focusing on the classificatory question at issue here, in chapters 3 and 4.

[23] *Aesthetics*, p. 138.

[24] Beardsley also allows that the sentence may be "indirectly self-contradictory" because certain presuppositions embodied in the predicative term may be violated by the subject (*Aesthetics*, p. 142). Notice, however, that as stated this account can deal only with utterances in which the predicate is metaphorical. Many metaphors do not fit this mold.

[25] "The Metaphorical Twist," p. 264.

[26] "The Metaphorical Twist," p. 271.

[27] *Aesthetics*, p. 143.

usually multiply ambiguous to a high degree. The sentence's larger context then eliminates, by pragmatic means, any irrelevant or contextually inappropriate connotations from the final interpretation.[28]

The interesting question is how we are to understand ‘connotation’. In order for it to be plausible that connotations are *semantic* features — features of "the significations of the words themselves,"[29] as Beardsley repeatedly and heatedly insists — connoted characteristics need to be relatively highly conventionalized. And this fits well with the way in which Beardsley initially introduces the notion. When a speaker uses a term, we can on that very basis assume that she probably also believes that the entities referred to have certain further characteristics; these further characteristics, Beardsley says, are the term's connotations.[30] This definition is sufficiently inclusive as to allow for the sorts of features that raised trouble in the examples cited above, such as (11) and (12). However, once again, even this definition proves to be insufficient to account for the full range of metaphorical interpretation.

In order to deliver the intuitive meanings for all metaphors — especially but not exclusively for rich poetic metaphors — and to explain the way in which metaphor "augments the resources of our language,"[31] Beardsley is forced to allow that connotations can also include highly unconventional features. These ‘unnoticed’ features, Beardsley says, "may wait, so to speak, lurking in the nature of things, for actualization — wait to be captured by the word... as part of its meaning in some future context."[32] Beardsley provides the following example to make this point. In interpreting the metaphorical phrase ‘th’ inconstant moon’, he supposes that we initially come up with no possible connotations that could serve as the metaphorical meaning. At that point,

---

[28] Thus, Beardsley writes that "The metaphor is full and rich, apart from any context; indeed, the function of the context is rather to eliminate possible meanings than to supply them" (*Aesthetics,* p. 138). Prior to his discussion of metaphor *per se*, however, Beardsley distinguishes what a word connotes in general from what it connotes *in a particular context*, which is a subset of its potential range of connotation (*Aesthetics*, p. 125). The difference between the two positions — that the metaphorical term by itself denotes *all* of its connotations, and that it denotes only the contextually appropriate ones — is important because they produce very different word meanings, given that Beardsley's theory is a semantic one.

[29] " The Metaphorical Twist," p. 269.

[30] *Aesthetics*, p. 125. See also Black's description of the ‘system of associated commonplaces’ linked to a word: "Metaphor," p. 40.

[31] *Aesthetics*, p. 144.

[32] "The Metaphorical Twist," p. 271.

we seize upon the verbal opposition, all right, but when we look for relevant connotations we are balked.... And so we look about among the accidental or contingent properties of inconstant people in general, and attribute these properties, or as many of them as we can, to the moon. And these properties would, for the moment at least, become part of the meaning of 'inconstant,' though previously they were only properties of those people. Then we might say that the metaphor transforms a *property* (actual or attributed) into a *sense*.[33]

The idea is that those 'accidental or contingent' features have now been made available for future metaphorical use. Repeated use, Beardsley continues, may eventually transform these features from connotations into defining features of the term's denotation. This does seem to be right. Many of our words do seem now to have literal meanings with apparently metaphorical roots: for example, (river) 'mouth', (table) 'leg', 'text'.[34] But if metaphor is granted the power to transform properties into senses in this way, then Beardsley's insistence on offering a semantic account of metaphor has been undermined. Far from depending on the predicate's established connotation for its meaning, as he initially claimed, metaphor turns out to *create* that connotation, by exploiting features of the objects that the word refers to. It thus seems that we are inevitably forced to include more than semantic meaning in our account of metaphorical meaning. We always need to appeal in at least some cases to our more general knowledge of how the world is.

As we saw with the lexicalists, we would have a nice argument for the distinctively semantic status of metaphor if metaphor *required* semantic incongruity, and if that incongruity, plus other features of word meaning, *sufficed* to deliver metaphorical meaning. These are the initially plausible necessary and sufficient conditions on which Beardsley builds his theory. But once again, and for much the same reasons, neither of these claims can be made good — even on Beardsley's quite inclusive notion of word meaning.

None of this shows that the initial intuition behind semantic theories of metaphor must itself be explained away. Metaphorical meaning *is*, after all, intimately bound up with the literal meanings and connotations of the uttered sentence's constituent terms. An adequate theory needs to explain this. Eva

---

[33] "The Metaphorical Twist, p. 273, emphasis in original.
[34] Though 'potted histories' of these metaphorical etymologies are notoriously inaccurate.

Kittay[35] in particular has <mark>explored the systematic ways in which metaphorical interpretation depends not just on a term's definition and connotations, but also on its surrounding 'semantic field'.</mark> This is the collection of syntagmatic relations (relations specifying the grammatical role of a term), contrast sets (for instance, 'summer' is a member of the cyclically ordered contrast set <seasons: spring, summer, autumn, winter>) and affinity sets (sets ordered by hyponymy and the part-of relation, among others) into which a given term enters. Perhaps all of this data is sufficiently conventionalized to be considered semantic; I return to this issue in §5.2. I am suspicious of this claim: I am inclined to treat much of it as non-semantic common background knowledge about the world and about typical word use. <mark>But even if we grant this point, we have also seen that metaphorical meaning often also depends essentially on information that is decidedly *not* part of a term's meaning, even when 'meaning' is understood in the relatively broad senses given by Cohen and Margalit, by Beardsley, or by Kittay.</mark> Thus, the truth behind the intuition that metaphor is intimately bound up with word meaning is not sufficient to establish the conclusion that metaphor is a semantic phenomenon.

1.1.2: *The New Semanticists*

Recently, some philosophers have begun to articulate another set of intuitions in favor of semantic theories, and to develop semantic accounts of a new sort based on these intuitions.[36] <mark>The leading idea here is that metaphorical utterances play a direct, explicit, and systematic role in conversation, as typical pragmatic phenomena do not.</mark> More specifically, a metaphorical utterance serves as the vehicle for making an explicit speech act, and thereby systematically alters the effects that later

---

[35] *Metaphor: Its Cognitive Force and Linguistic Structure*. See also Andrew Goatly, *The Language of Metaphors*, and Patti Nogales, *Metaphorically Speaking*.

[36] Proponents of this sort of view include most of the Relevance theorists and those closely associated with them: see especially Dan Sperber and Deirdre Wilson, *Relevance: A Theory of Communication and Cognition,* "Loose Talk," and "Rhetoric and Relevance"; Robyn Carston, "The Semantics/Pragmatics Distinction: A View from Relevance Theory"; and François Recanati, "Literal/Nonliteral." David Hills ("Aptness and Truth in Verbal Metaphor") defends a semantic account of metaphor of the sort outlined, but also maintains that metaphor works to affect content in a distinctive way: it involves 'make-believe' in the sense discussed by Kendall Walton (*Mimesis as Make-Believe)*. Josef Stern (*Metaphor in Context*) defends a contextualist semantic view of metaphor, but he maintains that context intrudes into semantics for metaphor in a more systematic and rule-governed way than other contextualists think it does.

uses of the same words have. Anne Bezuidenhout offers the following exchange to demonstrate these explicit, systematic effects:

> (13) A: How about Bill?
> B: Bill's a bulldozer.
> A: That's true. But isn't that a good thing in this case? We want someone who'll stand up to the administration and get things done for our department.
> C: I disagree that he's a bulldozer; that exterior hides someone who's basically insecure. But either way, Bill wouldn't make a good chair.

Here, the conversants offer their own judgments about the content that the speaker expressed metaphorically. Bezuidenhout says that "the most natural construal to put on such a dialogue is that B says something, that B says it with assertoric force, and that A and C are either agreeing or disagreeing with what B says."[37] By contrast, none of this is true of B's response in the following exchange:

> (14) A: How would you assess Mr. X's philosophical ability?
> B: Mr. X's command of English is excellent, and his attendance at tutorials has been regular.[38]

Here it is not nearly so natural to say that B *said* that Mr. X is not a good philosopher, although this is what B implicates. Nor is it normally natural or appropriate for A to respond to B's claim by agreeing or disagreeing with the implicature itself. The facts that metaphorical utterances but not implicatures allow their speakers to make explicit speech acts with contents distinct from their literal meanings, and that later hearers can take up those same metaphorical terms for making their own speech acts, might seem to suggest, then, that in metaphor the *words* uttered have been given new roles to play in that context of utterance. Another way to put this claim is to say that the words themselves take on a new, if temporary, semantics.

Rather than building the resources necessary for metaphorical interpretation into the standing meanings of terms in a language, as the 'old-school' semanticists did, the 'new' semanticists advocate a

---

[37] "Metaphor and What is Said: A Defense of a Direct Expression View of Metaphor," p. 157. Bezuidenhout employs this argument to show that metaphor affects *what is said*, and so should be treated as a form of direct expression. She explicitly refrains from claiming that metaphor a *semantic* phenomenon. This also seems to be the position of Sperber and Wilson. I largely ignore this distinction in what follows, because it would take us too far into delicate issues about how best to understand the phrase 'what is said'. I think the considerations of the next two sections, and especially the discussion of Hills's discourse embedding argument, tell against taking metaphor to be a form of direct expression, whether that is thought of as semantic or pragmatic. See §6.1.1 for more on the strict and the loose uses of 'says'.

[38] Grice, "Logic and Conversation," p. 33.

brand of *contextualism* about utterance meaning across the board. They argue that if we restrict our

semantics to just those phenomena that are fully conventional, or that at most depend on context in

heavily rule-bound ways (such as the contextual dependence of 'I'), then we will be left with a deeply

impoverished semantics. The 'meanings' of sentences delivered by the semantics will be mere abstract

theoretical constructions, with little or no relation to ordinary language users' intuitions about the content

and truth of utterances. As a result, we will have no genuine empirical constraints on our semantic

theorizing; and our theory will be of little use in explaining actual linguistic processing. In the place of

such a restricted, rigid and artificial semantic theory, these theorists think, we should adopt a

thoroughgoing semantic contextualism. On this view, all sorts of pragmatic factors can enrich and alter

the content that is directly — that is, semantically — expressed by an utterance. Metaphor should be seen

as just one more such factor.[39]

1.1.2.1: Psychological Realism

As I've just said, one prime impetus behind 'direct expression' views is a desire to reflect the

psychological reality of interpretation: to tie the theory of meaning as closely as possible to language

users' intuitions about both the process of comprehension and the intuitive truth-conditions of utterances.

François Recanati applies this motivation to metaphor in his article "Literal/Nonliteral." If we attend to

our ordinary intuitions about language use, he claims, then we will see that metaphor is in an important

sense a form of direct or *literal* expression, however paradoxical this may initially sound.

Recanati first sets to work on clarifying what it means for an utterance to be "nonliteral in the

ordinary sense." He says that there must be "something special about that use *that is, or can be,*

perceived by the language users themselves."[40] That is, the utterance's 'specialness' or nonliterality must

be "transparent" to ordinary speakers. This sort of nonliterality must be distinguished in principle, he

---

[39] See Mark Crimmins, "Hesperus and Phosphorus: Sense, Pretense, and Reference," for the application of similar
contextualist views to the semantics of belief reports; see also Keith DeRose, "Assertion, Knowledge and Context"
for a contextualist view of the semantics of knowledge ascriptions.
[40] "Literal/Nonliteral," p. 270, emphasis in original.

thinks, from the more theoretical notion of indirect or "*secondary* meaning — meaning derived from some more basic, primary meaning that it presupposes."[41] When language users arrive *directly* at an intended interpretation, and treat it as the *primary* meaning of the sentence, then the utterance is 'p-literal' ('p' standing for 'primary'); secondary meaning is p-nonliteral.

With these specifications in hand, Recanati then claims that although metaphor is the paradigm case of nonliterality 'in the ordinary sense', nonetheless "in its most central varieties" it should be counted as 'p-literal': it is primary in the sense that it comes first in the order of interpretation. To support this claim, he argues that an ordinary hearer "*readily* understands what is said by" an utterance of

(15) The ATM swallowed my credit card.

The hearer "*unreflectively* constructs" the relevant sense, he says, "*without going through a two-step procedure* involving the prior computation of the 'literal' meaning (whatever that may be) and a secondary inference to the actual meaning."[42] He thinks that typical conversational implicatures and irony, by contrast, do not work in this way. So, an utterance of

(16) Paul looks thirsty,

meaning that Paul might like a drink, or of

(17) Paul is a fine friend,

meaning that he is anything but, do require a 'two-step computational procedure'. They therefore count as *both* p-nonliteral and as nonliteral 'in the ordinary sense'.

Recanati tries to blunt the extreme oddity of his claim that metaphor is a literal use of language by pointing out that it is not an isolated phenomenon. Utterances' intuitive truth-conditions often depart from their conventionally encoded semantic content even though the utterances themselves are still 'p-literal'. The model case of this phenomenon is 'meaning enrichment'. So, for instance, an utterance of

(18) Have you had lunch?

is typically understood as asking whether the hearer has eaten lunch *today*, and not just at any time in the past. Similarly, an utterance of

---

[41] "Literal/Nonliteral," p. 266.

(19) Jack and Jill got married and had a child

will usually be understood to claim that the events occurred in the order mentioned. In these cases, Recanati maintains, what the speaker is intuitively taken to have *said* — not just to have communicated in some way or other — is a richer and more specific proposition than that encoded by the sentence's conventional meaning. Metaphor should be seen as just another instance of this widespread, systematic 'intrusion' by pragmatics into semantics. Metaphor differs from enrichment, he claims, primarily in that with metaphor the alteration consists in shifting or *extending* the semantically encoded content, so that it now literally applies to situations that it conventionally would not. By contrast, 'enrichment' typically adds additional specifications to the encoded content, as in (18) and (19), and thereby *narrows* the expressed content.[43]

Does this general argument show that metaphor is a case of direct expression or 'p-literality'? Recanati's justification for his claim that metaphors are p-literal crucially hinges on applying his 'transparency condition' to the criterion of p-*non*literality itself:

> p-nonliteralness is transparent to the language users... (This transparency is not a contingent property of p-nonliteralness. It is definitive of p-nonliteralness that the sort of inference at issue is conducted at the 'personal', rather than sub-personal, level and is therefore available to the language users.)[44]

But Recanati then vacillates in what he takes such 'transparency' to require. In this passage he assumes that it means that language users must be aware of going through an *inferential process* of deriving the

---

[42] "Literal/Nonliteral," p. 271, emphasis added.

[43] In the same vein, Anne Bezuidenhout takes metaphor to be a case of "loosening" the content expressed from the content strictly semantically encoded, where enrichment narrows the content expressed. ("Metaphor and What is Said: A Defense of a Direct Expression View of Metaphor," p. 168.) This view of metaphor as 'loosening' or 'extending' meaning traces back to Sperber and Wilson's article "Loose Talk," in which they argue that metaphor simply falls further out along the continuum of departure from conventionally encoded content than the more straightforward cases of looseness. A typical example of the latter would be using 'flat' to describe a surface which is largely planar but contains many nicks and bumps when viewed up close. Again, Bezuidenhout, like Sperber and Wilson, retain a distinction between semantically encoded, literal content and non-literal meaning, and therefore treat metaphor as a case of non-literal speech, albeit as one of 'direct expression'. Because of this, I do not focus on their views here.

[44] "Literal/Nonliteral," pp. 270-1.

secondary meaning. At other times, though, he takes it to mean only that "language users are aware of the distinction between the two layers of meaning as well as of the connection between them."[45]

The former requirement on p-nonliterality is untenably strong in general, and not just in application to metaphor. Relying on my own introspective powers, I'm not conscious of going through a two-stage inferential process to determine the relevant interpretation of utterances which Recanati does want to count as p-nonliteral, such as (16) and (17). And as Recanati himself admits, repeated use of a specific form of indirection, or explicit mention of contextual assumptions, will "short-circuit or conventionalize" what would otherwise be a two-step interpretive process into a single stage "— a complication," he says, "that I shall ignore."[46]

Thus, only the weaker condition for p-nonliterality can be warranted. And something like this condition *does* seem to capture the sense in which implicature and sarcasm are 'indirect' or 'nonliteral'. That is, I think that (16) and (17) above count as p-nonliteral because I — like ordinary, non-theorist speakers — think that there is something 'special' about the uses to which they are being put. I recognize that my intuitive sense of the utterance's intended interpretation differs markedly from its conventionally encoded semantic content, and I think that a two-stage process is required in order to explain how I could have determined the former on the basis of the latter.

However, Recanati invokes the stronger condition in defending his claim that metaphors are p-literal: as we saw above, he argues that ordinary hearers "readily" and "unreflectively" grasp the intended sense, without being aware of any inferential derivation. And the stronger condition is obviously essential to his position. Without it, most metaphors will count as p-nonliteral after all, because they do clearly meet the weaker condition on non-literality. That is, ordinary speakers in general *are* manifestly "aware of the distinction between the two layers of meaning as well as of the connection between them" when it comes to metaphor.

---

[45] "Literal/Nonliteral", p. 270. The first quote follows immediately upon the second, and is intended to be a restatement of it.
[46] "Literal/Nonliteral," p. 270.

Dead and dying metaphors like (15) provide the strongest cases for Recanati, but they are not examples of "metaphor in its most central varieties,"[47] as Recanati claims. Metaphorical utterances range across a spectrum of how radically they depart from conventionally encoded content, and of how much interpretive effort they require. (15) does seem like a case of relatively routine meaning 'extension'. George Lakoff and his colleagues[48] have demonstrated the pervasiveness of cases like this, which we are likely to interpret unreflectively in the course of conversation, but to count as nonliteral upon reflection or interrogation. At the one extreme of the spectrum, we may be unsure whether utterances like

(19) I finally got my ideas *across* to him

should be counted as metaphorical at all, or simply as literal uses that exploit new conventionalized meanings with metaphorical roots. At the other extreme, poetic metaphors like

(20) Juliet is the sun

do not depend on meaning 'extension', but on something more like wholesale *transfer*. Such metaphors often flaunt their literal absurdity, and require (and invite) sustained interpretive effort. In between, many ordinary conversational metaphors, such as

(21) Sam is a pig

or

(22) Bill is a bulldozer,

are on a par with most ironic utterances and implicatures in terms of how radically they depart from encoded content, how strong our ordinary intuitions of literal falsity are, and how much interpretive effort they require.

Recanati admits that metaphor ranges across a spectrum in this way. Strangely, though, he wants to count even poetic metaphors — metaphors in which the intended interpretation departs dramatically from the encoded semantic content — as p-literal, albeit as what he calls *figurative* p-literal utterances.[49] I don't know what this could even mean; I suppose the idea is that the metaphorical meaning still plays a

---

[47] "Literal/Nonliteral," p. 271.
[48] In *Metaphors We Live By, More than Cool Reason*, and other works.
[49] "Literal/Nonliteral," p. 272.

primary role in the order of interpretation even though it is 'transparent' to ordinary speakers that this meaning is 'special'. But even if we do find some way to make sense of 'figurative p-literality', his claim that metaphors can be 'figurative p-literal utterances' undercuts his original argument for the p-literality of metaphor, *even* under the unsustainably strong condition for p-nonliterality. These metaphorical utterances are often not primary in the order of interpretation, and do require sustained, conscious interpretive effort. But now we want to know what all these instances of metaphor — 'p-literal' or not, 'figurative' or not — have in common: what makes us treat them all as manifestations of the same phenomenon.

So it turns out that far from respecting language users' ordinary intuitions about meaning, as it set out to do, a 'direct expression' theory of metaphor like Recanati's forces us to disregard those intuitions. Ordinary speaker's intuitions about metaphorical meaning are that it "derives from a more basic, primary meaning which it presupposes."[50] And for metaphor, as for implicature, ordinary language users are prepared at least to try to reconstruct the interpretive process by which they moved from the sentence's conventionally encoded meaning to the intended meaning. By contrast, neither of these points apply with nearly the same force for utterances with 'meaning enrichment'. Thus, respecting the non-literality of metaphor need not rule out simultaneously accepting Recanati's arguments for the 'p-literality' of 'enriched' utterances.[51]

There may still be independent facts about how metaphorical utterances behave in conversational context which motivate treating both metaphorical and conventionally encoded literal meanings as semantic. But if these arguments can indeed establish that metaphor is semantic, then the semantics will at least need to be a two-leveled one. That is, it will still need to recognize the distinction between literal and metaphorical meaning if it wants to respect ordinary speakers' linguistic intuitions. This is the sort of semantic theory offered by Eva Kittay, in *Metaphor: Its Cognitive Force and Linguistic*

---

[50] "Literal/Nonliteral," p. 270.
[51] Although speakers' intuitions about the 'directness' or 'literality' of such utterances are not as universal or as stable as Recanati makes them out to be, either. See Kent Bach, "Seemingly Semantic Intuitions."

Structure. In the next section, I consider some arguments for this position: that metaphorical meaning should indeed be counted as semantic, along with literal meaning.

1.1.2.2: 'Lodging in the words'

Recanati's apparently bizarre willingness to countenance poetic metaphors as both p-literal and figurative suggests that a second motivation for 'p-literality', which he mentions only in passing, may actually be more fundamental to his thinking here. This is that "the situation talked about is 'categorized' by means of the words that are used to describe it."[52] Other writers on metaphor have shared this intuition: that in speaking metaphorically, speakers don't simply employ sentences as one-off tools for communicating some distinct content, as seems to happen with typical implicatures. Rather, such speakers use their words to present a broader conceptualization or 'schematization' of the subject being talked about. This schematization employs many of the distinctions embodied in the literal meanings of the words uttered, but redeploys them to carve up a new domain.[53] This fact about 'schematization' is then supposed to explain the systematic effects that metaphorical utterances have on later uses of the same and related words. Recall that these systematic effects were nicely exemplified in Bezuidenhout's 'bulldozer' exchange:

> (13) A: How about Bill?
> B: Bill's a bulldozer.
> A: That's true. But isn't that a good thing in this case? We want someone who'll stand up to the administration and get things done for our department.
> C: I disagree that he's a bulldozer; that exterior hides someone who's basically insecure. But either way, Bill wouldn't make a good chair.

In C's statement, 'bulldozer' contributes the same metaphorical content as in B's initial utterance.

Similarly, if C responded to B by saying

> (23) Really, Bill's just a cyclone fence,

---

[52] "Literal/Nonliteral," p. 271.

[53] Recanati quotes Langacker's discussion of schematization in "Foundations of Cognitive Grammar I"; see "Literal/Nonliteral," p. 271. Bezuidenhout expresses what I take to be the same point with her talk of an *ad hoc* concept; see "Metaphor and What is Said," pp. 157, 168. See also Eva Kittay, *Metaphor*; and Lynne Tirrell, "Extending: The Structure of Metaphor."

then that utterance would continue to deploy the same general range of assumptions about construction

sites that were activated by B's initial utterance, even though a new word, and with it a new concept,

would be invoked.

The leading idea is that speaker's application of a complex conceptual schema to a new domain

"lodges in the words"[54] she utters, so that those words, and others related to them, take on the job of

expressing those new categories in that conversational context. However, for the fact that a speaker has a

broader 'schematization' in mind to show that metaphor should be treated as a distinctively semantic

phenomenon, it's not enough that the speaker systematically exploit many of the distinctions embodied in

the literal meanings of her words, or even that her hearers also exploit those same distinctions in their use

of the same words. It must also be true that the *words* she utters now have the job of marking those

distinctions and applying their new content to the domain being talked about. And this claim requires

further argumentation. Otherwise, we still have the option of explaining the phenomenon of sustained

metaphoricity as a relatively systematic but still temporary case of speakers' meaning departing from

semantically encoded meaning.

David Hills and Josef Stern have each produced arguments intended to establish precisely this

further claim: that the words themselves have taken on a new, contextually dependent meaning. They

both argue that metaphors exhibit distinctively semantic behavior when embedded into larger contexts.

Hills focuses on embedding into larger stretches of discourse, while Stern focuses on embedding into

larger sentences.

Hills begins with an example much like Bezuidenhout's (13), but focused on a more poetic

metaphor. Thus, he points out that if someone responds to Romeo's utterance of

(20) Juliet is the sun

by saying "No, she isn't" or "She sure is," then this would naturally be construed as a response to the

metaphorical content of the initial utterance: to the claim that Juliet possesses those particular qualities of

---

[54] David Hills, "Aptness and Truth in Verbal Metaphor," p. 127. Hills here echoes Davidson's more sceptical use of the phrase in "What Metaphors Mean."

being gloriously, supremely beautiful, essential to his life, nurturing of his soul, and so on. Hills then

contrasts this with an utterance of

(14) Mr. X's command of English is excellent, and his attendance at tutorials has been regular.

In this case, the same sort of response ("Yes, he is" or "No, he isn't") will most naturally be taken to

accept or challenge what the speaker *said* — the claim that Mr. X's command of English is excellent and

his tutorials attendance regular — rather than the implicature that there's nothing positive that the speaker

can say to recommend him *qua* philosopher. From this difference, Hills concludes:

> So it would appear that Romeo's meaning gets lodged in Romeo's words in a way that Grice's meaning (in the letter of recommendation example) never gets lodged in Grice's words. The words of Romeo's utterance, as used by him on a particular occasion, get taken so as to express a thought they wouldn't express if they were taken literally — one which may be true or false or indeterminate in its truth value, one to which we are free to respond in ways that are appropriate only to thoughts that speakers have actually put into words.[55]

The argument seems to be that because other, later uses of the words 'bulldozer' and 'sun' inherit the

metaphorical content from the initial utterance, and so can be used to respond to that metaphorical

content, therefore we should think that the metaphorical meaning is a feature of the words themselves in

that conversational context, and not just of each speaker's use of them.

Now, it is true that someone can respond to Romeo's utterance of (20) by agreeing or disagreeing

with the content metaphorically expressed, for instance by uttering:

(24) She sure is.

Here the utterance is elliptical for

(24') She sure is the sun,

and 'sun' inherits the metaphorical interpretation of Romeo's initial utterance. One can also disagree

with the claim made, by saying something like

(25) No she's not. She's just a pallid, pasty little girl. Rosaline is the real sun around here.

This response would accept Romeo's intended use of 'the sun', but reject his particular application of it to

Juliet.

---

[55] "Aptness and Truth in Verbal Metaphor," p. 127.

However, these are not the only available responses. Someone can also reject the truth of Romeo's utterance construed *literally*, by saying

(26) Juliet isn't the sun; it's up there, and she's just a normal girl, here on Earth.

It is true that a hearer who was informed, charitable, and cooperative would likely not respond in this way, because it's contextually obvious that Romeo didn't intend his words literally. But precisely by speaking metaphorically, Romeo has made himself vulnerable to a response like (26). For instance, while he might feasibly object to (26) by saying "That's not what I meant," he cannot appropriately object by saying "That's not what I said." His hearer can always insist on a literal, strict interpretation of his words, and Romeo must grant that he wasn't meeting this established standard: he didn't mean that Juliet is in fact *the sun*. That he didn't really mean what he said is exemplified by the fact that he would not be willing to grant obvious and immediate consequences of this claim, such as that Juliet is up in the sky very far from earth, that she is composed of hot gas, that she is older than the earth, and so on.

Notice further that yet another sort of response to Romeo's utterance is also possible. One can use the very same words, and use them metaphorically, but not to echo the content of Romeo's intended claim, but rather to express a very different content. So, for instance, Benvolio might respond to (20) with something like:

(27) Yes, she sure is: just as the sun blinds and burns if you come too near, so will Juliet, a Capulet, burn you if you pursue her. Just as the sun is far off in the celestial distance, so is Juliet unreachable by you, a Montague. Indeed, just as the sun toils through its fixed rounds every day, so is Juliet trapped in the monotonous circle of Veronese social life. Better to focus on that ever-changing, ever-fascinating moon: Rosaline.

Now, if the initial metaphorical utterance of (20) had genuinely 'lodged' a new meaning in the *words* uttered, then that meaning should necessarily be inherited by any later use of those same words in that same context which responds to the initial claim. Any *other* way of using those words should be as irrelevant to the initial utterance as using 'bank' in its riverside sense would be in response to an utterance that employs 'bank' in the financial institution sense. This is so precisely because on Hills's hypothesis, the metaphorical utterance has generated a new meaning for the words uttered, and so if there are now two available meanings, they must be independent and separate, much as the two meanings of 'bank' are.

But, as we've seen, the speaker's metaphorical meaning is not necessarily inherited in the way that a genuinely semantic meaning would require. First, both the literal meaning and Romeo's metaphorical meaning are available as candidate contents for Benvolio's response. Therefore, Romeo can't have *said* something that has a single, metaphorical, meaning as its content. Second, and more importantly, Benvolio's own later use of those same words can serve to express either the literal content, Romeo's intended metaphorical content, or else the content of yet another claim; and any of these uses will themselves count as legitimate responses to Romeo's initial utterance. Therefore, Romeo's utterance can't really have fixed a new, non-literal meaning for his own use of the relevant words in the conversational context.

If we turn now to Stern's argument for metaphoric 'lodging', we find it too relies on a point about the behavior of metaphor in larger contexts. Here the claim is that the interpretation of whole sentences (as opposed to whole discourses) is constrained by the metaphorical interpretations of their constituent parts in a distinctively semantic way. To support this claim, Stern cites the following sentences as "semantically ill-formed" examples of verb phrase anaphora:

(28) The largest blob of gases in the solar system is the sun, and Juliet/Achilles *is, too.*
(29) Juliet is the sun, and Achilles *is, too.*

The only available interpretations of the second conjuncts in these cases are supposed to be those determined by a charitable interpretation of the first conjunct. This is supposed to be precisely analogous to the way in which, for

(30) John may leave tomorrow, and Harry, too,

the interpretation of 'may' in the second conjunct is automatically 'copied' from its interpretation in the first conjunct, so that if we read the first conjunct as saying that John *has permission* to leave, we simply cannot read the second one as saying that it's *possible* for Harry to leave. Thus, for our initial 'sun' sentences, the claim is that 'is, too' in (28) must be interpreted literally, while the same phrase in (29) must be interpreted to apply those features that are metaphorically predicated of Juliet. Because these interpretations of the second conjuncts are absurd, the sentences are supposed to be ill-formed. And

because, Stern says, "it is difficult to see how we might account for these constraints [whose violation produces the ill-formedness] in terms of use or mutual beliefs and expectations," therefore "we must attribute to metaphor the semantic structure, or meaning, necessary for the requisite condition to apply."[56]

Stern's own theory explains the constraint on metaphorical interpretation in the following way. He postulates an unpronounced operator [Mthat] at the level of logical form, which is analogous to David Kaplan's [Dthat] operator[57] but functions to make the term or phrase to which it is appended contextually sensitive in ways specific to metaphor. He argues that (28) is ill-formed because at the level of logical form, 'is the sun' in the first conjunct does not have an [Mthat] operator, and so it forces the elliptical 'is too' in the second conjunct to be read without one as well. (29) is less infelicitous than (28), on his view, because 'is the sun' in the first conjunct of (29) *is* prefixed by an [Mthat] operator, and so the 'is too' of the second conjunct inherits a metaphorical reading. Stern puts this claim by saying that both [Mthat 'is the sun'] and the elliptical 'is too' have the same metaphorical *character*, in Kaplan's sense of that term. But (29) still counts as semantically ill-formed, he thinks, because the properties which form the *content* of 'is the sun' in the first conjunct are inappropriate for the subject to whom the elliptical 'is too' is predicated in the second conjunct.

My first point in response is that I simply disagree with Stern on the data. I think it is indeed possible to read (28) and (29) so that each conjunct is interpreted separately. Indeed, to my ears (29) is not even that infelicitous, although it's true that a speaker would need to go on to explain her utterance. Stern himself admits all this, up to a point; he merely claims that these "imposed interpretations" carry "a feeling of play or pun," and insists that this feeling reveals the semantic ill-formedness in question.[58] However, we don't yet have an independent argument that the "feeling of play or pun" reveals a distinctively *semantic* infelicity. And the fact that such interpretations are available for (28) and (29),

---

[56] *Metaphor in Context*, pp. 70-1.
[57] Cf. Kaplan, "Dthat."
[58] I also reject the claim that *if* the second conjunct of (29) *is* read literally, then it must be nonsensical. See my "The Generality Constraint, Nonsense, and Categorial Restrictions."

however odd they may be, essentially distinguishes them from a genuinely semantic constraint of the kind manifested in (30).

Indeed, much of the same oddness is preserved if we split the original conjuncts up into separate sentences:

(28') The largest blob of gases in the solar system is the sun. And Juliet/Achilles is the sun.
(29') Juliet is the sun. And Achilles is the sun.

But if this is so, then the infelicity cannot be the result of violating a specifically semantic constraint. In these revised sentences there is no longer any ellipsis that could force the same reading for both terms.[59] Rather, I think, the infelicity of both the original conjoined sentences and of their separated twins stems from the fact that the interpretation of the first verb leads us to expect a similar interpretation for the second. This expectation makes (28) and (29) more or less ineffective and potentially misleading ways to express yourself if you intend the conjuncts to be interpreted separately, because they make salient certain interpretive assumptions which must then be ignored.[60]

Finally, notice that the sense of a distinctively *semantic* ill-formedness could result entirely from the inappropriate use of a single definite description to refer to two different things in the same context. If we replace (28) and (29) with

(28'') The largest jewel in the world is a diamond, and Juliet/Achilles *is, too.*
(29'') Juliet is a diamond, and Achilles *is, too*,

then the infelicity does seem to be diminished significantly — though again, not eliminated fully. Notice further that on Stern's view, a sentence like

(31) Juliet and Achilles are each diamonds in their own very different ways

should be uninterpretable, because it requires binding over different metaphorical readings of 'is a diamond'. To my ear, though, a metaphorical interpretation of (31) is fine and felicitious.

---

[59] Cf. Zwicky and Saddock, "Ambiguity Tests and How to Fail Them."
[60] Bezuidenhout suggests that the unacceptability of (28) arises because of an equivocation in the use of 'is'. She argues that the first conjunct employs the 'is' of identity while the second conjunct employs the 'is' of predication, on the grounds that the first, but not the second, is reversible. The fact that (29) does not equivocate in its uses of 'is' would then explain its greater acceptability ("Metaphor and What is Said," p. 181).

Bezuidenhout, Hills and Stern have pointed to a pair of genuine facts about metaphor. **First**, a metaphorical utterance makes a speech act in which the speaker's intended content is available for explicit response. **Second**, after an initial metaphorical utterance, it is quite easy for hearers to employ the same manner of speaking to express their own attitudes toward the content communicated. Those ways of using those words and other, systematically related words may even have become the default uses in that context. These facts do distinguish metaphor, at least in degree, from many other indirect speech acts. They therefore call for an explanation; and it does not seem that a pragmatic account which simply assimilates metaphor to the general case of implicature will succeed in explaining them. However, these facts do not themselves establish that the metaphorical meaning has really 'lodged in' the words themselves, and so is genuinely semantic meaning. I offer my own explanation of these phenomena in §5.4.5, after my own account of metaphorical meaning is on the table.

It would be a relief if we did not need to treat metaphor as a semantic phenomenon, because doing so would tax the semantics enormously. As we've seen, there are two basic routes for incorporating the resources needed for metaphorical interpretation within a semantic theory. **One** can either enrich words' conventional meanings, or **one** can make word meaning contextually dependent. The resources needed to interpret metaphor are so rich and broad that including them all within words' conventional meanings would effectively explode the semantics altogether. Some form of contextualism is thus the only viable semantic option. And mainstream theories do already admit a significant degree of contextualism into semantics. Demonstratives, tense operators like 'was', and implicit comparatives like 'small' all require appealing to context to fix their content. Further, the form of contextual dependency they exhibit is not nearly as fixed and algorithmic as, for instance, the rule for interpreting 'I' is. In addition to such indexical elements, we must also appeal to context to disambiguate the speaker's intended meaning from among alternative readings of sentences like

(32) Visiting royalty can be boring.
(33) I visited the bank today.

However, the sort of context-dependency that metaphor displays is importantly different from those of either indexicality or ambiguity. With metaphor, the context-dependency cannot always be traced to particular syntactic elements within the sentence, let alone to elements that we have independent reason to think require contextual specification as part of their conventional meaning, as indexicals do.[61] Nearly any word or phrase can be used metaphorically, and can be used to communicate a host of different contents. Because of this, if we decide to treat the contextual sensitivity of metaphor as semantic, then we will have to build this contextual sensitivity individually into the semantic rule for nearly every word in the language, and the semantic rule for contextual dependence will in each case need to be so general as to be nearly vacuous.

By contrast, if we put metaphor to one side for a moment, then the sorts of words which independently require contextual supplementation are much more restricted, and the sorts of contextual supplementation they call for are much more clear-cut. This means that treating metaphor as a context-dependent semantic phenomenon would put a much greater strain on a compositional analysis than the semantics otherwise needs to bear. Indeed, it would effectively force us to abandon any hope of developing a relatively systematic, explanatory compositional semantics. But **compositionality** seems to be the only viable explanation for the massive generativity of our linguistic production and comprehension. We may be forced to abandon the hope for a compositional semantics for other reasons — because of polysemy or 'loose talk', for instance. But we should prefer a more conservative explanation of metaphor, one which does not rule out the very possibility of such a semantics, if one is available.

---

[61] See Jeffrey C. King and Jason Stanley, "Semantics, Pragmatics, and the Role of Semantic Content."

1.2: *Pragmatic Theories of Metaphor*

Pragmatic theories of metaphor[62] start from the intuition that in speaking metaphorically, we undertake speech acts with propositional contents distinct from that fixed by the conventional meanings of the sentences uttered. So, for instance, an utterance of

(34) I will trek through this treacherous minefield with you

might be used make a promise: that the speaker will accompany the hearer through a contextually salient dangerous situation, say a legal case. And an utterance of

(35) Go out there floating like a butterfly and stinging like a bee

can be obeyed without the addressee becoming airborne: he just needs to step lightly and jab quickly and aggressively.

According to the standard pragmatist, a metaphorical utterance is *indirect*, in the sense that the speaker's meaning is not directly expressed by what she actually says. Pragmatic theories typically explain the connection between what is said and what is meant by postulating that in speaking metaphorically, speakers intend for their hearers to recognize the following three things:

(a) that the speaker is unlikely to be expressing the proposition, *P*, that is fixed by the sentence's conventional meaning,[63] because doing so would conflict with mutually shared assumptions about the conversation to this point and about the world more generally;
(b) that the speaker is therefore speaking metaphorically;
(c) that the most likely interpretation of the speaker's intended meaning is therefore *q*, given further mutually shared assumptions, made relevant by the particular sentence uttered.

As I suggested in the last section, one *prima facie* source of evidence for the claim that metaphor is a function of the speaker's intentions in this way is the specific form of context-dependency that metaphor manifests. The context-dependence of metaphorical utterances often cannot be traced back to any independently-motivated context-dependence that is produced by the conventional meanings of particular constituents within the sentence uttered. Instead, both whether the utterance counts as

---

[62] See Grice, "Logic and Conversation"; John Searle, "Metaphor"; Bach and Harnish, *Linguistic Communication and Speech Acts*; Merrie Bergmann, "Metaphorical Assertions." Emma Borg's view ("An Expedition Abroad: Metaphor, Thought, and Reporting") is in most respects a pragmatic one, although it does not rely on a Gricean pragmatic mechanism of the sort described below.

metaphorical at all, and which particular content it expresses, seem to depend on the speaker's manifested communicative intentions and attitudes. So, for example, an utterance of

(36) George is the next Napoleon

 might be used to claim that George will overthrow dissipated authority, that he is a megalomaniac despot, merely that he is a short up-and-coming politician — or something else altogether. Which of these is expressed will depend on whether the speaker intends to make a laudatory or a disparaging claim, and more specifically on the speaker's manifested attitudes toward the bearers of the names mentioned. This evidence for a pragmatic theory is no more than *prima facie*: so far, one could insist on treating the sort of context-dependency metaphor exhibits as semantic. Perhaps the claim here amounts to not much more than the claim that metaphor has the 'feel' of a pragmatic phenomenon. But, as I noted at the end of the last section, the fact remains that metaphor's context-dependency is sufficiently pervasive and messy to make a pragmatic analysis attractive.

We also find support for the pragmatic schema above in our actual interpretive practices. It is true, as Recanati says, that hearers often do not engage explicitly in the postulated reasoning processes. And we usually cannot provide fully explicit, deductively valid justifications of our metaphorical interpretations. But it is also true that when a metaphor is particularly obscure or puzzling, then our patterns of thought do approximate the schema above. For both routine and puzzling metaphors, we do take speakers to have committed themselves to claims and other speech acts by their metaphorical utterances. And we do require there to be some sort of reasonable connection between the sentence uttered and the speech act made, as (c) suggests. When speakers or hearers are challenged to explain how they could have meant, or have taken someone to have meant, *that* by saying *this*, we do expect them to be able to provide some sort of reconstruction, however rough, of the connection between the two. (By contrast, we do not ordinarily expect this sort of reconstruction for cases of semantic context-dependency.) And if a challenged metaphorical interpretation cannot be justified by some sort of rational

---

[63] Recall that I include within 'conventional meaning' appeals to context which are required by the conventional meanings of the words uttered, even if this appeal to context includes an appeal to the speaker's intentions in that

reconstruction — if our interlocutor simply asserts that that's what the utterance means — then we will reject it.

Semanticists sometimes object to a pragmatic theory of metaphor by saying that it puts things the wrong way round: we deduce a speaker's communicative intentions from the meanings of her words, they argue, and not vice versa. But the only meanings that words have are semantic meanings; therefore, they conclude, a theory of metaphor must be a semantic theory, a theory of metaphorical word meaning. Thus, Beardsley writes:

> We do not decide that a word in a poem is used metaphorically because we know what the poet was thinking; rather we know what he was thinking because we see that the word is used metaphorically. The clues to this fact must be in the poem itself, or we should seldom be able to read poetry.[64]

Given the pragmatic schema above, though, we can see why this is not a sound objection. We do not need first to know *what* a person is thinking before we can recognize *that* an utterance is metaphorical. Rather, the hearer at least implicitly reasons through (a) and (b) before arriving at (c). We do determine that an utterance is metaphorical (step [b]) by relying on clues available within the poem or the conversation to this point (step [a]). But this does not itself show that these 'clues' to metaphoricity must be semantic. Indeed, even semanticists like Beardsley and Kittay acknowledge that the relevant clues are often pragmatic: the literal meaning need not be logically absurd, but is still inappropriate because it is not the sort of thing that someone would be likely to claim, at least not in this context.[65] Once we have determined that the utterance is metaphorical, then we seek out clues to what the speaker might mean. One absolutely essential set of clues here is again provided by the literal meanings of the words uttered and their mode of combination. But again, we also rely on the larger conversational or poetic context,

---

context.

[64] "The Metaphorical Twist," p. 269. (Beardsley also claims here that to describe an utterance as one in which the speaker says something without meaning it is thereby to treat that utterance as a case of irony. This is absurd.) Cf. also Kittay: "We do not need to rely on any actual intentions for meaningfulness; rather what we do is impute intentions, and the imputed intentions derive from the regular and systematic interrelations of terms within a language" (*Metaphor*, p. 48; see also p. 114).

[65] Kittay makes a valiant, and interesting, attempt to reduce this infelicity to semantic deviance by extending semantics to include a more general, heavily enriched form of discourse representation theory (*Metaphor,* chapter 2). However, even if we are willing to grant her the claim that discourse representation of the sort she describes *is* semantic, still even she hesitates to assume that all the indications of 'deviance' that trigger metaphorical interpretation will count as semantic.

and on everything we know about the world and the speaker's specific psychology. This sort of reasoning — combining semantic with worldly knowledge — is typical of pragmatic processing generally.

So we have some general motivations for treating metaphor along the general pragmatic lines sketched above. However, the mere plausibility of this sort of schema obviously can't establish it as an adequate theory of metaphor. Indeed, it suffers from some obvious, serious problems, centering around the 'therefore's in (b) and (c). First, how is the hearer supposed to get from the probable non-literality of the utterance to the conclusion that the speaker intends her utterance in a specifically metaphorical manner — as opposed to, say, an ironic one? Second, and more seriously, how is the hearer supposed to determine what the speaker's intended metaphorical meaning *is*? That is, what are the distinctive, mutually shared patterns and processes of thought on which metaphorical comprehension relies, and in which the speaker intends her hearer to engage in order to recover that metaphorical meaning? Without answers to these questions, the pragmatist has merely motivated the claim that metaphor fits the general profile of indirect speech. We have been given no analysis of *how* metaphor works, or indeed of what metaphor *is*.

No one has a particularly satisfying answer to the first question, of how a specifically metaphorical interpretation is indicated. We can say something, although nothing very surprising, about the basis for determining that a speaker does not intend her utterance literally ([a] in the schema above): if interpreted literally, the speaker's utterance would violate the Gricean Cooperative Principle and at least one of its corollary Maxims.[66] But underwriting the inference from (a) to (b) is more important, and more difficult.[67] There are no obvious distinctive syntactic, semantic or pragmatic markers for metaphoricity, as opposed to non-literality and indirectness more generally. The consensus seems to be that a

---

[66] See Grice, "Logic and Conversation"; also Bach and Harnish, *Linguistic Communication and Speech Acts*, ch. 4.
[67] We should also note that sometimes the speaker indicates (b) explicitly, for instance by prefacing her utterance with a phrase like 'Speaking metaphorically...'.

distinctively metaphorical interpretation just 'leaps to mind', and that the speaker chose her words knowing that this would be a likely interpretation.[68]

Perhaps, though, the consensus opinion can provide us with an answer to our question after all. Faced with a metaphorical interpretation which could plausibly express the sort of thing the speaker might have in mind, perhaps we can appeal to the easy availability of this interpretation as itself a reason for thinking that the speaker intended her words metaphorically in the first place. Given this possible avenue of explanation — and given that no other answer seems to be forthcoming — I will focus on the second question posed above: what are the distinctive principles of metaphorical comprehension? That is, what sorts of shared assumptions are relevant? And what is it about those assumptions, and about the role that they play, that makes an interpretation metaphorical?

No existing pragmatic account delivers satisfactory, substantive answers to these questions, though, either.[69] John Searle is in advance of most pragmatic theorists in actually attempting to provide such principles for metaphorical interpretation. He accepts as his task the criteria we have just laid out: to identify a specific "shared system of principles" on the basis of which we can "provide a rational reconstruction of the inference patterns that underlie our ability to understand" metaphors.[70] He also agrees that we should state these principles "in a way that makes it clear how metaphorical utterances

---

[68] See e.g. Sperber and Wilson, "Pragmatics, Modularity, and Mind-Reading." Sperber and Wilson think that their Principle of Relevance — communication maximizes cognitive effect while minimizing cognitive effort — guarantees that the first interpretation a hearer entertains will be the intended one. Part of the question in which we're interested, however, is precisely how speakers and hearers manage to coordinate expectations in this way, in general and in the specifically metaphorical case. For this, their invocation of the Principle of Relevance does not provide a satisfying answer: it simply assumes that there *is* an answer.

[69] Grice is typical in regard. He treats metaphor as a form of conversational implicature, and requires that for conversational implicatures in general, "The presence of a conversational implicature must be capable of being worked out; for even if it can in fact be intuitively grasped, unless the intuition is replaceable by an argument, the implicature (if present at all) will not count as a conversational implicature; it will be a conventional implicature." (p. 31). But then all he says about the way the implicature is 'worked out' for metaphor is that "The most likely supposition is that the speaker is attributing to his audience some feature or features in respect of which the audience resembles (more or less fancifully) the mentioned substance." This is extremely vague, but few later theorists have come up with anything much more specific. In particular, the invocation by Relevance theorists of 'loosening' constraints encoded in a literal interpretation, and of thoughts more or less tightly 'resembling' their semantically encoded content is just as general. (See e.g., Sperber and Wilson, "Loose Talk", Carston, "The Semantics/Pragmatics Distinction"; Bezuidenhout, "Metaphor and What is Said", and Robyn Langdon, Martin Davies and Max Coltheart, "Understanding Minds and Understanding Communicated Meanings in Schizophrenia".)

[70] "Metaphor," p. 104.

differ from other sorts of utterances in which speaker meaning does not coincide with literal meaning."[71]

His failure to fulfill his announced criteria is thus particularly instructive.

Searle focuses on the simple case in which the speaker says that *S* is *P* and means that *S* is *R*. He treats my (a), the determination of non-literality, and (b), the determination of metaphoricity, in the schema above as a single step. He breaks my final step (c), the identification of the metaphorical meaning, down into two subsidiary interpretive stages. These are (1) the determination of a range of *possible* metaphorical meanings, and (2) the selection of the *appropriate*, intended metaphorical meaning from among these possibilities. He offers eight principles which govern (c1), and which presumably serve to meet the criteria for a theory of metaphor above. His principles for deriving '*S* is *R*' from '*S* is *P*' are as follows. (I have introduced the notation 'means$_{met}$' to abbreviate 'can be uttered metaphorically to mean that'; I have also systematized the presentation of the examples, especially of their paraphrase.)

1. Things which are *P* are by definition *R*; usually *R* will be one of the salient defining characteristics of *S*. Example: "Sam is giant" means$_{met}$ "Sam is big."
2. Things which are *P* are contingently *R*; again, *R* will usually be a salient or well-known property of *P* things. Example: "Sam is a pig" means$_{met}$ "Sam is filthy, gluttonous, and sloppy, etc."
3. Things which are *P* are often said or believed to be *R*, even though both speaker and hearer may know that *R* does not in fact apply to *P* things. Example: "Richard is a gorilla" means$_{met}$ "Richard is mean, nasty, prone to violence, and so on."
4. It is a fact about our sensibility, whether culturally or naturally determined, that we just do perceive a connection, so that *P* is associated in our minds with *R*. Example: "Sally is a block of ice" means$_{met}$ "Sally is unemotional."
5. The condition of being *P* is like the condition of being *R*. Example: "You have become an aristocrat" means$_{met}$ "Your new status is like that of being an aristocrat."
6. *P* and *R* are the same or similar in meaning, but one, usually *P*, is restricted in its application, and does not literally apply to *S*. Example: "His brain is addled" (no interpretation provided).
7. A principle extending the simple '*S* is *P*' form to other syntactical forms, basically by applying 1-6 at a higher order. Example: "The ship ploughs the sea" means$_{met}$ "The ship moves the sea to the side of the prow as it moves forward."
8. *P* and *R* may be related as part-whole or container-contained, so that metonymy and synecdoche also count as metaphors.

What are we to make of this? Searle seems to be right that the interpretation in each case depends on something like the principle cited. However, we have no indication so far of what unites them: of why they are all principles for interpreting metaphor *per se*, and why these are the only such principles. We had hoped, as Searle himself said, that by isolating these principles we could identify some systematic

---

[71] "Metaphor," p. 78.

and distinctive feature(s) of metaphorical interpretation. And something like this does seem to be possible for sarcasm and for hyperbole. In the former case, the speaker means to claim something like the 'opposite' of what she actually said; while in the latter, the speaker means to claim that a less extreme version of the situation described actually obtains.[72] But so far, neither necessary nor sufficient general features of this sort are forthcoming for metaphor.

In addition, each of the principles adduced is itself so broad, and the list as a whole comprises so many different ways in which P and R might be related, that in the end they amount to not much more than the requirement that P and R must be similar (or just related) in some respect or other.   And as philosophers enjoy pointing out, similarity or relation without further specification is a trivial relation that holds among any set of objects. Indeed, on Searle's account, in the following conversation B's utterance meets the criteria for being a metaphorical utterance:

> (37) A: How many degrees do you think Bill's got?
> B: Well, he's a professor.
> A: So he must have a Ph.D., right?
> B: Exactly.

Being a professor is contingently but commonly and saliently associated with having a Ph.D. Thus, when B says that Bill is a professor and thereby implicates that he has a Ph.D., Searle's principle 2 for metaphorical interpretation is satisfied. If *this* is all there is to the "shared system of principles" that underlies metaphorical comprehension, then it looks like every utterance should be equally interpretable as a metaphor, and in indefinitely many ways. But clearly this is not all there is, given the success we often do display at agreeing which utterances are metaphorical, and at converging on common metaphorical understandings.

The real challenge to explaining metaphorical interpretation, then, is not to set out principles for *generating* a list of features that the uttered predicate P and the intuitively intended predicate R have in

---

[72] See, e.g. Merrie Bergmann, "Metaphorical Assertions"; Bach and Harnish, p. 67. This minimal account of sarcasm is not adequate without an account of what function 'the opposite' is, and this seems to vary across cases of sarcasm (see Stern, *Metaphor in Context*, pp. 232-238; Fogelin, *Figuratively Speaking*, pp. 9-10). Many theorists also seem to assume that a theory of sarcasm is a theory of irony, but there are many cases of irony which do not fit this 'opposite' function view. Much irony seems to work by mocking, in a more complex way, the situation as it

common (stage [c1]). Rather, the difficulty is to provide principles for *narrowing* the indefinitely long list of features that could be relevant to those that actually are relevant in any given case (stage [c2]). But about this step, Searle says only that we should look back to the subject being discussed: "The basic principle of this step is that only those possible values of *R* *which determine possible properties of S* can be actual values of *R*."[73] This is some help, but "possible properties of *S*" is still much too general, and still fails to indicate anything specifically distinctive of metaphor. Many, perhaps most, of his candidate *R* values will be logically, metaphysically and physically possible properties of *S*.[74] Having a Ph.D. is certainly a possible property of a professor. We still need a principle for selecting the most plausible metaphorical values for *R*. I suspect that Searle has in mind some sort of pragmatic possibility, along the lines of 'properties that the speaker could possibly intend to predicate of *S*, given the conversational context.' If so, then by itself, this criterion is still much too weak: without further specification, many *P*-values are likely to be contextually plausible *R*-values. And in any case, this criterion does not advance us far beyond the initial schema (a) though (c) above, with its general, vague appeal to relevance and 'mutual assumptions.'

Searle responds to my example (37) in the following way.[75] The fact that the principles he adduces apply not just to metaphor, but to indirect speech and implicature more generally, should not be seen an objection to his account but as a further advantage of it, he says. It demonstrates that we are appealing to principles of interpretation we already need for other purposes. The real reason that B's utterance can't count as metaphorical, he argues, is that there is nothing wrong with it at a literal level. By contrast, if someone had uttered

(38) My dog is a professor,

---

would be represented by someone who actually meant what the speaker says (see Sperber and Wilson, "Irony and the Use-Mention Distinction").

[73] "Metaphor," p. 112, emphasis added.

[74] Notice that speaker may also want to make claims which are metaphysically impossible but epistemically possible. These are presumably *not* possible properties of *S*.

[75] personal communication.

then this utterance could be metaphorical, because it can't be literally true. Thus, the real work of determining *whether* an utterance is metaphorical occurs not at his second stage of interpretation ([c1] in my schema) but at the first (my [a] and [b]). *After* a determination has been made *that* the utterance is metaphorical, then his principles come in to provide possible metaphorical meanings.

I think that Searle's response to this objection puts his account in a worse position, not a better one. By shifting the burden of responsibility onto his first interpretive step in this way, he seems to have abandoned his best resources for explaining metaphor "in a way that makes it clear how metaphorical utterances differ from other sorts of utterances in which speaker meaning does not coincide with literal meaning,"[76] as he set out to do. As Searle himself points out in his article, and as the authors I cited in §1.1.1 have also shown, it is not criterial of metaphor that the sentence uttered be literally absurd, as (38) is. Searle's own example is Disraeli's utterance of

   (39) I have climbed to the top of the greasy pole.

It is not even necessary that the sentence be pragmatically inappropriate, as Cohen's example,

   (10) Moscow is a cold city,

shows.

Further, even if we accept that there must be some sort of deviance in order to trigger the search for a non-literal interpretation, as there often in fact is, still Searle has offered no basis for identifying the speaker's intention to be interpreted in a specifically metaphorical way (step [b] in my schema above). And as I said above, ==I don't think there *is* any criterion — syntactic, semantic, or pragmatic — on the basis of which a hearer could identify an utterance as metaphorical in advance of actually interpreting it.== In particular, most other types of implicature and figurative language also often (though again, not always) rely on some sort of literal deviance to trigger the search for an alternative interpretation. Searle defines 'metaphor' relatively broadly, so as to include synecdoche and metonymy, but surely he does not want to define it so broadly as to include irony and implicature as well.

---

[76] "Metaphor," p. 78.

I think the best hope for explaining what is distinctive about metaphor is to focus on precisely the stage Searle does focus on in his article: (c1), identifying a "shared system of principles" which could explain how speakers and hearers manage to generate possible metaphorical interpretations. If we can identify a distinctive metaphorical connection between what is said and what is meant, of the sort his principles aim to provide, then we will at least have a substantive account of what metaphor *is*. And perhaps we can then use this connection to explain why hearers interpret utterances metaphorically in the first place, by arguing that the easy availability of a plausible, relevant metaphorical interpretation is itself a reason to think that the speaker intended her utterance metaphorically.

It's not as if Searle has gone off track in his account in some obvious way — either in general, or in his principles specifically. Each of his interpretive stages is necessary and important. Each of his principles for (c1) is relevant to metaphorical interpretation, and it's not obvious how to subsume them under a more general principle that still substantive. My own suspicion is that the real problem lies in seeking inferential principles that would render fully explicit the obviously enthymematic arguments that we intuitively offer in trying to cash out the vague principle (c) in particular cases. The rational reconstructions that we use to make sense of communication generally and of indirect speech more specifically are never tight deductive arguments; they always rely on somewhat vague appeals to what is contextually salient.[77] However, even by the relatively relaxed standards usually set for such arguments, the reconstructions for metaphorical interpretation that we offer in particular cases look more like 'abductive miracles' than like inferences from reasonable, independently salient assumptions.[78] The fact that Searle's analysis ends up being so vague is a natural consequence of this. What we need, I think is an

---

[77] Cf. Bach and Harnish, pp. 80-81:
"Detailed as it is, the SAS [Speech Act Schema] does not represent the precise form of inference (to be) made by the hearer. Left open are the questions of just which mutual contextual beliefs are activated and of just how they enable the hearer to find the right candidate for the speaker's illocutionary intent (not that there is ever any guarantee of success). A complete account of the hearer's inference would require a systematic theory of saliency...To our knowledge no such theory has yet been developed. Until it has, the SAS can represent only the general pattern of the hearer's inference."

[78] I believe I first heard the phrase 'abductive miracle' from David Hills.

alternative way of connecting and moving between thoughts that is at once less specific and more substantive than the sort of 'system of inferential principles' that Searle offers.

1.3: *Non-Cognitivist Theories of metaphor*

The failure to provide substantive principles that explain how metaphor works vitiates existing pragmatic theories of metaphor. But these theories also typically suffer from another failure, this time one of neglect. They fail to acknowledge, let alone explain, the intuition that metaphors also work to bring about psychological changes that cannot be adequately captured in propositional or truth-conditional terms. Metaphors, especially rich, poetic metaphors, are important and powerful communicative tools precisely because they can induce in their hearers new ways of thinking and feeling about the subject under discussion. So, for instance, Romeo's utterance of

(40) But soft, what light through yonder window breaks? It is the east, and Juliet is the sun![79]

evokes a complex, nuanced set of attitudes toward Juliet. Similarly, Lear's description of his daughter Goneril casts both her, and familial relations more generally, in a distinctive, highly negative light:

(41) But yet thou art my flesh, my blood, my daughter — Or rather, a disease that's in my flesh, which I must needs call mine. Thou art a bile, a plague-sore, or embossed carbuncle, in my corrupted blood.[80]

An adequate account of how metaphor works must explain these rich psychological effects as well; pragmatic theories, as they stand, do not.

Pragmatic theorists have tended to miss this fact because they have focused on, and so have dealt best with, metaphors as they occur in everyday conversation. Metaphors like

(21) Sam is a pig

and

(42) Sally is a block of ice

usually serve to communicate just one or a few claims about their subject; their non-propositional effects are comparatively minimal. It is no mere accident, though, that pragmatic theorists have focused

---

[79] *Romeo and Juliet*, II.ii.1.
[80] *King Lear,* II.iv.242-246.

primarily on such cases. Standard philosophical theories of language lack the explanatory resources they would need to make sense of communicative effects other than the transmission of truth-conditional contents more generally, not just in the case of metaphor.

Many mainstream theorists excuse this deficiency by maintaining that such effects pertain merely to the psychology and sociology of language use, and so lie beyond the purview of a theory of meaning.[81] However, this position distorts the philosophy of language in general, and it is untenable when it comes to the investigation of metaphor. If we *systematically* use language to do something other than transmit propositional content, then philosophers need to explain how we do this. At a minimum, if certain words are conventionally used to communicate certain affective attitudes, then this use is neither the result of a mere idiosyncratic causal association, nor even just a generalization about the causal effects which the word produces in an entire population. This sort of use of words is of a piece with the standard intentional communicative structure outlined by Grice, Searle, and others: a speaker uses her words so as to produce in her hearer the effect of recognizing the effect which she is trying to produce in him. In this case, the effect in question is the recognition that the speaker has a certain affective attitude, and perhaps that she intends for her hearer to share that attitude as well.[82]

Turning to metaphor specifically, an adequate theory must explain what metaphor distinctively accomplishes, *qua* metaphor, across the full range of cases. We will not have answered that question unless we address metaphor's salient non-propositional effects. Perhaps not all metaphors work in just the same way, or accomplish just the same effects. But we should not begin by erecting an artificial boundary within the class of metaphors at the outset, and simply assuming that the non-propositional effects which metaphors accomplish are of no concern to a philosophical theory of metaphor. We will only be in a position to sort out the relation between metaphor's propositional and non-propositional functions after we have conducted a close examination of the phenomenon of metaphor as a whole.

---

[81] See e.g., Kent Bach, "The Semantics/Pragmatics Distinction: What It Is and Why It Matters," p. 67.
[82] Cf. Robert Brandom, *Making It Explicit*; Jennifer Hornsby, "Meaning and Uselessness: How to Think about Derogatory Words"; Kaplan, "The Meaning of 'Ouch' and 'Oops'."

Further, by focusing only on metaphor's role in the communication of propositional content —
content which can often also be stated directly and literally — we run the risk of making metaphorical
speech out to be an inefficient detour employed only for the sake of rhetorical ornamentation or flourish.
We then need to explain why speakers so pervasively indulge in such behavior, even in serious contexts
such as scientific theorizing, and even — perhaps especially — in urgent, practically-oriented
communicative contexts such as political debate. Something must be wrong with any theory that makes
such a common use of language appear to be strangely convoluted and marginal, as pragmatic theories
too often do.

Non-cognitivist theorists of metaphor, exemplified by Donald Davidson,[83] do focus on precisely
the non-truth-conditional effects that are ignored by mainstream accounts, but they share the mainstream
accounts' assumption that only truth-conditional communication can be relevant to a theory of meaning
per se. Non-cognitivist theories have both a negative and a positive component. On the negative side,
they are 'non-cognitivist' in the sense that they deny that metaphorical utterances have any distinctive
propositional, or cognitive, content. As Davidson says, "metaphors mean what their words, in their most
literal interpretation, mean, and nothing more."[84] Crucially, they deny further that the *speaker* manages to
mean anything else, "no matter how indirectly," by speaking metaphorically.[85] In order to establish these
points, non-cognitivists typically challenge defenders of metaphorical content to specify that content
precisely for novel, poetic metaphors, like this line of Auden's:[86]

(43) The hourglass whispers to the lion's paw.

As Davidson says, "It should make us suspect the theory that it is so hard to decide...exactly what the
content is supposed to be."[87]

---

[83] See e.g. Donald Davidson, "What Metaphors Mean"; Richard Rorty, "Unfamiliar Noises I: Hesse and Davidson
on Metaphor"; Marga Reimer, "Davidson on Metaphor."
[84] "What Metaphors Mean," p. 245.
[85] "What Metaphors Mean," p. 259.
[86] Reimer, "Davidson on Metaphor," p. 142. I'm not sure that I *do* understand this metaphor, but my best shot at a
paraphrase (out of the context of the poem) is that the passing of time threatens the current incarnation of power.
[87] "What Metaphors Mean," p. 262.

Non-cognitivists admit that a hearer may sometimes be led to entertain thoughts *as a result of* hearing a metaphorical utterance, and that the speaker may even have intended for her utterance to lead the hearer to entertain those thoughts.[88] That is, they admit that certain cognitive effects may be included among metaphor's perlocutionary effects. Even then, though, they maintain that the relationship between the utterance and the thoughts evoked is not one of meaning or claiming anything. As Davidson says, "Joke or dream or metaphor can, like a picture or a bump on the head, make us appreciate some fact — but not by standing for, or expressing, the fact."[89]

Rather than being in the business of expressing content, non-cognitivists argue, the function of a metaphor is to lead us to "see the world in a new light," or under a new *aspect*.[90] Metaphors provide us with a kind of "lens or lattice" which, as Max Black says "selects, emphasizes, suppresses, and organizes features of the principal subject."[91] This reorganizational "filtering" then "nudges us into noting"[92] novel or surprising similarities among things. Because this aspectual effect is neither "finite in scope" nor "propositional in nature," attempting to cash it in for a paraphrase is "simply misguided."[93]

Where Davidson focuses on the way in which a state of 'seeing-as' leads us to notice hitherto neglected similarities, other theorists have concentrated on its power to arouse feelings and attitudes in its hearers. It is this power, they claim, that accounts for metaphor's special rhetorical force. 'Perspectives' and 'aspects' are more forceful, immediate, and pervasive than propositional thoughts. The fact that

---

[88] "What Metaphors Mean," p. 255.

[89] "What Metaphors Mean," p. 262. Note that this is a tendentious analogy because on many views, pictures and dreams do represent (putative) facts.

[90] **There are a number of tensions between Davidson's theory of metaphor in "What Metaphors Mean" and his more general philosophical views. One** important one is between his positive invocation of 'seeing-as' here and his claim elsewhere (e.g., "A Coherence Theory of Truth and Knowledge") that perception just *is* the causation of belief. **Another** is between his voluble insistence here that metaphors have no propositional content, and his invocation of speaker meaning in disputing Donnellan's treatment of the distinction between referential and attributive uses of descriptions: "Jones has said something true by using a sentence that is false. This is done intentionally all the time, for example in irony or metaphor." ("A Nice Derangement of Epitaphs," p. 470.)

[91] Max Black, "Metaphor," pp. 44-45. Black is the person Davidson is most concerned to criticize for his theory of metaphorical *content*, but he largely follows Black in what he takes to be metaphor's *effects*: "I have no quarrel with these descriptions of the effects of metaphor, only with the associated views as to *how* metaphor is supposed to produce them. What I deny is that metaphor does its work by having a special meaning, a specific cognitive content." ("What Metaphors Mean," p. 262.)

[92] "What Metaphors Mean," p. 253.

[93] "What Metaphors Mean," p. 263.

perspectives and aspects play such an important role in metaphor therefore explains why metaphors can have the powerful — even, according to Hobbes,[94] Locke[95] and Susan Sontag,[96] pernicious — effects that they do. As Richard Moran says, "Part of the dangerous power of a strong metaphor is its control over one's thinking at a level beneath that of deliberation or volition. In the mind of the hearer an image is produced that is not chosen or willed."[97] This image lingers, and continues to "frame" one's feelings about the subject characterized, even when one rejects any supposed content that the metaphor might convey. It is this sub-cognitive influence that then explains why, "if someone is described as having all the charm of a damp kitchen sponge, it's no good simply to *deny* it, after he or she has registered appreciation of the phrase."[98] The hearer's attempted denial is undermined by the metaphor's success: it has produced the intended image and feeling in him, and the very fact that he has gotten the image after all makes him "complicit" in the metaphor.[99] Thus, if, as Moran says, "the notions of communication and of saying require that a distinction can always be drawn between understanding and belief,"[100] then metaphor's 'complicity effect' — the fact that its effects are 'simply irresistible' once understood — would seem to show that metaphor is in something other than the business of communication.

==On the highlighted sentence==

In evaluating non-cognitivist theories, we need to distinguish between their negative and positive components. Non-cognitivists are right to point out that metaphor is importantly involved with making us

---

[94] "But for Metaphors, they are in this case utterly excluded. For seeing they openly professe deceit, to admit them into Councell, or Reasoning, were manifest folly" (*Leviathan*, p. 44).

[95] " We must allow that all...the artificial and figurative application of words eloquence hath invented, are for nothing else but to insinuate wrong ideas, move the passions, and thereby mislead the judgment; and so indeed are perfect cheats: and therefore...they are certainly...wholly to be avoided." *An Essay Concerning Human Understanding*; Bk. 3, Ch. 10 ("Of e Abuse of Words"), p. 34.

[96] Susan Sontag, *Illness as Metaphor and AIDS and Its Metaphors*.

[97] "Seeing and Believing: Metaphor, Image, Force," p. 90.

[98] "Seeing and Believing," p. 91.

[99] The idea of "complicity" here comes from Ted Cohen: "When the device is a hostile metaphor or a cruel joke requiring much background and effort to understand, it is all the more painful because the victim has been made a complicitor in his own demise" ("Metaphor and the Cultivation of Intimacy," p. 10). Wayne Booth formulates the point about metaphor's 'complicity' and 'irresistibility' thus:

"The speaker has performed a task by yoking what the hearer had not yoked before, and the hearer simply cannot resist joining him; they thus perform an identical dance step, and the metaphor accomplishes at least part of its work even if the hearer then draws back and says, 'I shouldn't have allowed that!'"

("Metaphor as Rhetoric: The Problem of Evaluation," p. 52).

'see' things under new aspects and feel about them in new ways. They are also correct that these psychological states are poor candidates for content: since they are not themselves propositional, they do not even represent anything truth-evaluable. But non-cognitivists often also assume that 'seeing-as' and feeling are the *only* available candidates for content.[101] And this is just not so, as the many examples I've already cited demonstrate. When non-cognitivists do take up the possibility that metaphors might have some propositional content after all, the only candidate that they seem to consider is a proposition to the effect that the things seen anew are *alike* in some way.[102] It is true that people have held views of this kind about metaphorical meaning — views on which metaphors are really elliptical similes or statements of comparison.[103] But once again, this is not the only available position. In particular, the 'seeing-as' effects of metaphors don't just 'nudge us into noting' *similarities* between things, but can also lead us to notice that the thing(s) under discussion possess certain inherent properties. The proposition that the thing possesses those properties is one obvious, and *prima facie* plausible, candidate for the metaphor's content.[104]

Let us, then, take this proposition as a candidate content for metaphorical meaning for a moment. In the last two sections, we have seen some good reasons for agreeing that this can't be the meaning of

---

[100] "Seeing and Believing," p. 99. Presumably this distinction is required in order to retain in turn the distinction between illocutionary or pragmatic effects and perlocutionary ones.

[101] In the same manner, Rorty thinks it is an advantage of Davidson's account that it "makes it possible to see...metaphors as causes of our ability to do lots of other things — e.g., be more sophisticated and interesting people, emancipate ourselves from tradition, transvaluate our values, gain or lose religious faith — without having to interpret these latter abilities as functions of increased cognitive activity." ("Unfamiliar Noises I," pp. 284-5.) This would be an advantage for Davidson's account only if these were the sorts of things that people had claimed to be the cognitive content or meaning of metaphors. I know of no one who has thought this. Black and Kittay, at least, do often seem to think that the 'new way of seeing' that a metaphor induces is its cognitive content, and so Davidson is right to point out that this is an untenable position. See e.g., Kittay, *Metaphor*, pp. 39 ff. and p. 96.

[102] Thus, Davidson writes: "A metaphor makes us attend to some likeness, often a novel or surprising likeness, between two or more things. This trite and true observation leads, or seems to lead, to a conclusion concerning the meaning of metaphors" ("What Metaphors Mean," p. 247). Davidson thinks we should keep the 'trite and true observation' and junk the fallacious conclusion. Similarly, Reimer asserts that someone speaking metaphorically "means to get the hearer to notice certain similarities between two things; she does not, *in addition*, intend to communicate any proposition to the effect that the similarities in question are there to be noticed." ("Davidson on Metaphor," p. 149, emphasis in original.)

[103] Robert Fogelin: *Figuratively Speaking*. In "The Metaphorical Twist," Beardsley argues that Paul Henle ("Metaphor") holds this view, but Henle does not in fact hold it. Searle and Black each attack this view as a common one, but cite no specific references.

[104] I'll argue in §5.4.3 that this is not in fact a satisfactory account, because it doesn't always deliver the intuitively appropriate content. But it is at least an initially feasible position.

the sentence uttered. But why can't a *speaker* use a metaphor to express or mean that a fact obtains?

Davidson himself is not as explicit about this as we might hope, but the reason seems to be that he doesn't

think the right sort of connection obtains between what is said and the thoughts it occasions. The sort of

connection that obtains with genuine meaning is rational and normative. In such a case, the speaker and

hearer are engaged in a joint, cooperative activity; the speaker exploits her knowledge of the hearer's

knowledge and reasoning processes to formulate her utterance so that, *if* her hearer is rational and

assumes that she is rational, *then* he *should* be able to determine what she intends to convey to him. By

contrast, Davidson assumes that the connection between utterance and thoughts occasioned in the case of

metaphor is a merely *causal* one, of 'nudging,'[105] 'prompting,'[106] 'inspiring,'[107] 'provoking, or

inviting.'[108] Because the connection here is merely causal, the idea runs, the thoughts occasioned by

metaphors can at most be part of the utterance's intended perlocutionary effect.

       Richard Rorty articulates what seems to be Davidson's position thus:

> In Quine's image, the realm of meaning is a relatively small 'cleared' area within the jungle of use, one whose boundaries are constantly being both extended and encroached upon. To say, as Davidson does, that 'metaphor belongs exclusively to the domain of use' is simply to say that, because metaphors (while still alive) are unparaphrasable, they fall outside the cleared area.[109]

Within the cleared area, Rorty goes on to say, linguistic behavior is "regular, predictable"[110]; here we can

legitimately speak of cognitive content, and with it of "norms and intentions,"[111] and "a place in a pattern

of justification of belief."[112] Outside of the cleared area, we can speak only of "stimuli,"[113] of "unfamiliar

noises"[114] which, like other "anomalous non-linguistic phenomena," "do not (literally) *tell* us anything,

but...do make us notice things... They do not have cognitive content, but they are responsible for a lot of

---

[105] "What Metaphors Mean," p. 253.
[106] "What Metaphors Mean," p. 263.
[107] "What Metaphors Mean," p. 257.
[108] "What Metaphors Mean," p. 261.
[109] "Unfamiliar Noises I," p. 285.
[110] "Unfamiliar Noises I," p. 285.
[111] "Unfamiliar Noises I," p. 292.
[112] "Unfamiliar Noises I," p. 295.
[113] "Unfamiliar Noises I," p. 291.
[114] "Unfamiliar Noises I," p. 293.

cognitions."[115] Metaphor should thus, on Rorty's view, be analogized to birdsong, thunderclaps, and sonatas: while each sort of noise has specific psychological effects upon us as a result of our past causal interactions with it, "this is not to say that the noises...have anything like non-natural meaning, even when these noises happen to be expressions of English, or notes on a musical scale."[116] If Rorty is right, then Wittgenstein's parody of the empiricist theory of ideas provides us with a true account of metaphor: "uttering a word is like striking a note on the keyboard of the imagination."[117]

In sum, the Davidson-Rorty line seems to be that meaning something, *P*, by one's utterance brings both the speaker and hearer into the rational, normative realm. It commits the speaker to defending certain further claims that are inferentially related to *P* if challenged; it also implies that the hearer's further thoughts can get that meaning right or wrong, and that there is a sense in which he should get it right.[118] The non-cognitivist claim is that metaphorical utterances do not fit this model.

But this claim is simply false. As we've seen, metaphors don't just 'prompt' and 'inspire' thoughts. They are also used to undertake speech acts with content distinct from the uttered sentence's conventional meaning. This is true not just of 'dead' metaphors like

(44) The city is located at the river's mouth

and

(45) John kicked the table leg,

nor just of relatively routinized ones, like

---

[115] "Unfamiliar Noises I," p. 290. Rorty's position can be summed up in the pithy slogan that follows this passage: "As with platypuses, so with metaphors."

[116] "Unfamiliar Noises I," p. 294, fn. 2.

[117] *Philosophical Investigations*, §6. The Wittgensteinian analogy nicely captures the causal nature of the connection between metaphor and its effects, as Davidson thinks of it. In elaborating Davidson's view, though, Rorty has attributed a **more radical** view to Davidson than Davidson himself holds. Davidson thinks that it's not just the causal history of our interactions with the uttered words that matters to its effects, but also our current understanding of the literal meaning of the uttered sentence. Davidson is perfectly willing to admit that such sentences have *some* non-natural meaning: their literal one. (This is a good thing, as we've seen, since sentences used metaphorically may be literally unexceptionable.) He is also willing to admit that our understanding of this meaning is relevant to the metaphor's effects; he just denies that the effects themselves can be counted as communicative content. Despite this rather serious distortion, I think Rorty does succeed in articulating the basic line of Davidson's thought about why even cognitive effects cannot count as speaker's meaning.

[118] This sort of commitment obtains even for speech acts with illocutionary forces other than assertion. If a speaker promises to make *P* the case, he also thereby promises to make some of the obvious logical entailments of *P* obtain

(22) Bill is a bulldozer

or

(46) Clifford is a mouse,

but also of full-blooded, rich, poetic metaphors, like

(20) Juliet is the sun.

==Hearers do take the speakers of such metaphors to have committed themselves to something which can be accepted or denied.== They do take themselves to have understood, or sometimes to have misunderstood, the speaker's claim. And they do offer additional evidence that confirms or contradicts the claim made. In serving to undertake speech acts in this way, metaphorical utterances of sentences class together with literal utterances and apart from non-sentential utterances.

Thus, what Rorty takes to be a strength of his account — that it assimilates metaphorical utterances of sentences to "scraps of poetry which send shivers down our spine, non-sentential phrases which reverberate endlessly, change our selves and our patterns of action, without ever coming to express belief or desires"[119] — instead betrays a fundamental blind spot. He is correct that from Yeats's metaphorical phrase,

(47) that dolphin-torn, that gong-tormented sea[120]

hearers cannot "acquire any beliefs which these particular words express",[121] but this is not because the phrase is just "unfamiliar noise," as he claims. ==Rather, the problem is the basic fact that (47) is not a sentence at all, and so not a suitable vehicle for staking any sort of claim.==

It is true, as non-cognitivists point out, that the content expressed by a metaphorical utterance may be ==open-ended==, unparaphrasable, and even partially indeterminate, especially when it comes to rich,

---

as well. If a speaker orders her hearer to make *P* obtain, she is thereby committed to the claim that *if* her hearer complies, then certain obvious logical entailments will also thereby obtain.

[119] "Unfamiliar Noises I," p. 285.

[120] The line is from "Byzantium."

[121] "Unfamiliar Noises I," p. 293.

poetic metaphors. ==But this does not distinguish metaphor in principle from other forms of speech.==[122] A speaker who says

(48) Jane is a real woman now

may mean something quite open-ended by her utterance, and not be able to express the full richness of her thought any other way. Likewise, there may be no semantic equivalent with which to paraphrase a demonstrative phrase like 'that sound'. Neither is metaphor unique in engendering non-cognitive effects. Laudatory and derogatory expressive terms do this as part of their conventional meaning; and particular literal utterances often do it pragmatically: an assertion counts as an insult, for instance, partly in virtue of its intended non-cognitive effects. Thus, there is no reason to choose between expressing cognitive content and causing us to 'see things in a new light' or feel a different way: metaphor can do both.

==The confusion about metaphor seems to have stemmed from the fact that Black and some other connotationists assumed that the 'filtering', 'seeing-as' effects of a metaphorical utterance just were its content. But once we accept something along the lines of an Austinian tripartite distinction between semantics, pragmatics, and perlocutionary effects, then it becomes obvious that not all of what a speaker intends to do with her utterance is necessarily part of the content she expresses with it.== Of course, as Richard Moran emphasizes,[123] if we do opt for a "two-aspect" theory of metaphor, on which it both communicates content and also 'frames' our way of thinking, then we will need to tell some story about the connection between these aspects. But this is an invitation for further work, not to throw our hands up in despair.

==A non-cognitivist theory of metaphor doesn't merely deny important facts about what metaphors are in fact used to do. Because it denies these facts, it ends up exiling metaphor from the sphere of philosophical explanation altogether.== This is a serious deficit, not just for proponents of the 'full-blooded' project of understanding how we use words to express thoughts about reality (a project which

---

[122] Cf. Kent Bach, "Speaking Loosely," p. 259; Emma Borg, "An Expedition Abroad: Metaphor, Thought, and Reporting," p. 244; Herman Cappelen and Ernie LePore, "On an Alleged Connection between Indirect Speech and the Theory of Meaning."
[123] "Seeing and Believing," p. 109.

Rorty condemns as "metaphysical"[124]), but for someone of a more minimalist, Davidsonian persuasion as well. Rorty makes much of the idea that Davidson has good reason to exile metaphors from the realm of meaning, because they are not "regular, predictable linguistic behavior."[125] Just as "physicists must simply disregard insensible perturbations and concentrate on relatively conspicuous and enduring regularities," he thinks, so too must the radical interpreter set aside metaphors, along with jokes and lies, as puzzling distractions.[126] But there are at least two essential disanalogies here. First, whereas the physicists' perturbations are by definition marginal data, metaphors and other nonliteral uses of language are pervasive to language use, both within and across speech communities. And, unlike with "insensible perturbations," the radical interpreter has no *a priori* way to distinguish nonliteral speech from the 'core' cases.

Second and more fundamentally, though, a fully developed, fully satisfactory physical theory *will* explain as much of the marginal data as possible. By contrast, Davidson and Rorty would have us exile today's metaphors from the realm of interpretation altogether,[127] not just set them aside as hard cases for later incorporation into the theory. Rorty's analogy is surely correct in suggesting that a radical interpreter should first get to work on isolating the most well-entrenched, systematic correlations between sentences uttered and their conditions of utterance. An interpreter will not be in a position even to distinguish speakers' lies, mistakes, jokes, malapropisms, and metaphors from sincere, truthful, literal discourse, let alone to distinguish among these various 'parasitic' uses, until he has developed a fairly robust, though provisional, interpretive theory. But surely the analogy is also correct in suggesting that once he is armed with such a provisional theory, the radical interpreter should get to work on explaining the remaining data. And it seems that a broadly Gricean mechanism like that offered in the pragmatic schema of §1.2 provides a smooth, minimally disruptive way to accomplish this.

---

[124] "Unfamiliar Noises I," p. 287.
[125] "Unfamiliar Noises I," p. 285.
[126] "Unfamiliar Noises I," p. 286.
[127] Rorty repeatedly emphasizes that the line between the metaphorical and the literal is continually shifting, but this concerns the possibility that *future* uses of what are now metaphors may become literal. It does not allow the

Finally, notice that by his own lights, Davidson is not really entitled to dismiss metaphor and other such uses of language as merely a matter for psychology, and so not appropriate to meaning proper. For, unlike Quine, Davidson thinks that an interpretive theory is essentially bound up with a theory of psychology: different interpretations of these puzzling utterances will determine different attributions of attitudes to their speakers, and vice versa. It is certainly possible that the data in question are just too complicated ever to systematize and bring within the theory of meaning. However, we should seek to explain them if we can. And given that ordinary language users evidently do have *some* sort of interpretive strategy that enables them to communicate metaphorically, it seems that a radical interpreter should be able to develop the necessary resources as well.

1.4: Seeing-as *and* Content

I have now provided sketches and assessments of the three main types of theories of metaphor. When we compare the pragmatic and non-cognitivist accounts as I've described them, we should notice that their strengths and weaknesses are strikingly complementary. The two theories are incompatible largely because they each focus exclusively on a different, but equally limited, range of cases. Pragmatists focus on ordinary conversational metaphors, and neglect the rich aspectual effects of their more poetic brethren. Non-cognitivists focus on just those metaphors and effects, but typically refuse to count ordinary conversational metaphors *as* metaphors because aspects play a less palpable role there.[128] We need to improve this dialectical situation by examining concrete examples from across the whole range of metaphor, and by teasing out the particular ways in which these examples are both similar to and different from one another.

More interestingly, pragmatic theories stand in need of a story about the distinctive psychological structures and processes that underlie metaphorical interpretation; only then will their suggestive schema have any real substance. Non-cognitivist theories do offer a psychological mechanism that they take to be

---

possibility that our theory may succeed in bringing what are at any given moment genuinely metaphorical utterances within the realm of meaning. The line shifts as the data shifts, not as the theory grows.

distinctive of metaphor, but they fail both to provide a substantive account of that mechanism, and to acknowledge the use to which language users actually put it.

If we could come up with a psychological mechanism of this sort, and show how it underwrites metaphorical interpretation, then we would have addressed the puzzles about metaphor with which we began, and in which theorists of metaphor have traditionally been most interested: How do speakers and hearers manage to converge on common metaphorical interpretations? What interpretive principles are specific to metaphor and distinguish it from other linguistic phenomena? Both semantic contextualists and non-cognitivists about metaphor have generally dismissed these questions as 'merely psychological', and hence non-philosophical (albeit perhaps interesting nonetheless). But if our task as philosophers of meaning is to elucidate the basic principles by which language users manage to communicate thoughts about the world, then there is no *a priori* reason to dismiss these questions for falling on the merely psychological side of the line.

In the light of this discussion, my own proposal for a theory of metaphor can be stated pretty simply. I agree with the non-cognitivists that 'seeing-as' is indeed crucial for understanding metaphors, both ordinary and poetic. But I also think that 'seeing-as' functions as the *means* rather than as the sole, final *end* of comprehension. Speakers intend for their hearers to recover their metaphorical content, I claim, by way of cultivating a state of 'seeing-as'; an utterance counts as metaphorical just in case the speaker intends that her hearer recover her communicative content in this way. By embedding an account of 'seeing-as' *within* an account of speakers' communicative intentions, we can do justice to the central insights of both the non-cognitivist and the pragmatist theories. We can both accept that non-truth-conditional effects are not good candidates for propositional content, and still explain how metaphors can have propositional contents. Of course, the fact that speakers put the effects of 'seeing-as' to further work does not mean that those effects are eradicated once they accomplish their intended task. Nor does it mean that speakers are not interested in those effects. Sometimes, producing these effects may even be

---

[128] Moran makes a similar point; "Seeing and Believing," p. 98.

their *primary* interest. Rather, it means that 'seeing-as' forms part, but not all, of the story about metaphor.

However, if my theory is to be more substantive than either the pragmatist or non-cognitivist accounts, then I need to say something more about just what it means to 'see something in a new light', or under a new aspect, and then explain how *this* enables hearers to recover speakers' intended meanings. The rest of the dissertation is devoted to this project.

CHAPTER 2: PERCEPTION AND SEEING-AS

      As we've seen, there is a strand of philosophical theorizing about metaphor, represented most

forcefully by Donald Davidson, which claims that metaphor "makes us *see one thing as another* by

making some literal statement that inspires or prompts the insight."[1] The idea is that a metaphorical

utterance both reflects and reinforces a certain 'perspective' on things for someone who already inhabits

it; and that it can be used to invite, lead, or induct a non-sharer into the same perspective. So, for

example, an utterance of

    (1) Man is a wolf

 is supposed to cause us to 'see' mankind as wolfish. An utterance of

    (2) Confusion now hath made his masterpiece!

is supposed to make us 'see' confusion as an artist intent on creating new and finely crafted artworks, and

as a result, given the context of utterance,[2] to 'see' Duncan's murder as a supreme act of chaos, one that

inaugurates a natural, national, and theological catastrophe. And Anthony's utterance of

    (3) You all do know this mantle...
       Look, in this place ran Cassius' dagger through;
       See what a rent the envious Casca made;
       Through this the well-beloved Brutus stabb'd,
       And as he pluck'd his cursed steel away:
       Mark how the blood of Caesar followed it,
       As rushing out of doors, to be resolv'd
       if Brutus so unkindly knock'd or no;
       For Brutus, as you know, was Caesar's angel.[3]

is supposed to make us 'see'... what? as what? Here the answer is not quite so clear. Are we to 'see'

Brutus' stabbing Caesar as the knock on a door? Or to 'see' the gushing blood as a page running out of

---

[1] "What Metaphors Mean," p. 253, emphasis added. Many people have shared the intuition that metaphor has
something important to do with changing our 'perspective' on the thing discussed; these include I. A. Richards,
William Epson, Paul Henle, Max Black, Richard Moran, and Eva Kittay.
[2] *Macbeth*, II.iii; Duncan's murder is just discovered (line 64):
    Macduff. O horror! horror! horror! Tongue nor heart
      Cannot conceive nor name thee!
   Macbeth. & Lennox. What's the matter?
   Macd. Confusion now hath made his master-piece!
     Most sacrilegious murder hath broke ope
     The Lord's anointed temple, and stole thence
     The life o' the building!
[3] Shakespeare, *Julius Caesar*, III.ii.170ff.

his master's house following a mysterious knock, as John Crowe Ransom suggests, and Monroe Beardsley vigorously denies?[4] Or something else altogether?

I want to leave aside for now the question of just what the relation must be between the sentence uttered and the way of 'seeing' it prompts. As the second and third examples here suggest, that relationship cannot always simply be read off from the words uttered and their syntactic mode of combination. (3) in particular suggests that sometimes ==a metaphorical utterance may reflect an underlying 'perspective' that is not explicitly indicated by the sentence uttered.== But we cannot even begin to deal with such complexity until we have a story about what a 'way of seeing' or 'perspective' or 'aspect' is in the first place.[5] In the next three chapters I develop such a story: an account of what phenomenon people are after when they invoke 'seeing-as' with respect to metaphor. Given how many people share the intuition that 'seeing-as' is important for metaphor, surprisingly few have attempted to articulate what this could actually amount to.[6]

2.1: *Seeing-As and Aspectual Thought*

The difficulty in getting a grip on what we're talking about here is compounded by the obvious but often neglected fact that we're not really concerned with actual perception.[7] ==For we typically encounter and employ metaphors in situations where the object or situation under discussion is not actually present== — for instance, when we are reading a book, or describing an experience to someone who has not shared it.

We might then think that *imagined* perception is the relevant phenomenon. Davidson and others have sometimes suggested that live metaphors work by encouraging us to conjure up and entertain images before our minds' eyes. In this vein, Davidson argues that the sentence

---

[4] Ransom, "Poetry: I. The Formal Analysis"; Beardsley, "The Metaphorical Twist."
[5] I return to the issue of how the sentence uttered determines the aspect generated in §5.2.
[6] This is true for discussions of 'seeing-as' in ethics as well. Philosophers like John McDowell ("Virtue and Reason") and Jay Wallace ("Addiction as a Defect of the Will") invoke the importance of *seeing* an action *as* virtuous, but don't explain what this involves.
[7] For others who make this point, see e.g., Neville Kemp, "Metaphor and Aspect-Perception," and Lynne Tirrell, "Seeing Metaphor as Seeing-as: Remarks on Davidson's Positive View of Metaphor."

(4) He was burned up

is no more than "a slangish idiom," on the grounds that it no longer makes us "picture fire in the eyes or smoke coming out of the ears."[8] His assumption here seems to be that *if* it were a live metaphor, then it *would* make us picture those things; presumably it did do this to its hearers when it was still alive. Davidsonian accounts of metaphor are sometimes called "image theories" for this reason.[9] The idea is that the utterance causes us to imagine seeing the situation described in the concrete terms of the predicating metaphor.

However, although metaphor does have an intimate connection with imagery, this can't provide an accurate analysis of the phenomenon, either. First, as Kendall Walton has pointed out,[10] we are especially drawn to employ concrete **metaphors in speaking about abstract objects**, like mankind, confusion, justice, and corporations. Being abstract, these latter things cannot even in principle be perceived or imaged by us. We cannot picture life itself as being any way at all when we think of it as "but a walking shadow" — although we may conjure up some image which is merely *associated* with it. Second, and more importantly, though, conjuring up images is just the wrong sort of activity in this context. If I, or even Romeo, spend too much time trying to *picture* Juliet as the sun, then we will at a minimum be distracted from, and at worst be defeated in, our understanding of the relevant metaphor. Indeed, we will be taking the metaphor too literally. David Hills makes a distinction in this context between two sorts of imagining: *fancying* and *modeling*.[11] When we fancy that one thing is something else, we do attempt to cultivate an imaginative experience of the one thing as the other: to picture how a friend looked when she was young, or how the Dean would look in a jester's costume, for instance. When we model one thing in terms of another, by contrast, we use one way of thinking to 'filter' another. Metaphor crucially involves only the latter notion.

---

[8] "What Metaphors Mean," p. 253.
[9] See Martin Davies, "Idiom and Metaphor."
[10] "Metaphor and Prop-oriented Make-Believe."
[11] "Aptness and Truth in Verbal Metaphor," p. 138.

Thus, the term 'seeing-as' applies not just loosely, but metaphorically, to metaphor. Nonetheless, the fact remains that in describing metaphors' effects, the language of vision — of 'seeing-as', 'seeing in a new light', and 'taking up a perspective' — is overwhelmingly natural. So we still need to pay attention to what it is about metaphor that makes these metaphors so appropriate for describing the non-perceptual effect that it induces. I'll call this non-perceptual phenomenon 'aspectual thought', as opposed to aspectual perception. The best way to work out the import of these visual metaphors for aspectual thought, I believe, is to take a close look at the perceptual phenomenon itself, and then to employ the insights gleaned there as guides for analyzing its non-perceptual analogue. To this end, in this chapter I investigate perceptual experience. I should emphasize at the outset, though, that I am not attempting to provide a full analysis of perceptual experience, or even just of seeing-as or aspectual perception. I am interested in these phenomena here only insofar as they provide us with clues for understanding the non-perceptional phenomenon of aspectual thought. Thus, my treatment of many important and interesting issues about perception will be quite cursory.

2.2: *Non-propositionality*

One important intuition behind invocations of 'seeing-as' in relation to metaphor is that the sort of understanding produced by metaphors is different in kind from that produced by most literal utterances. As Davidson says, because the effects of 'seeing-as' are "neither finite in scope" nor "propositional in nature," they are not the right sort of thing to be restated in literal, propositional form. An important first step, then, is to understand what it could mean to say that a psychological effect is 'non-propositional'.

In different contexts and for different purposes, theorists have used the term to mean quite different things. On one standard line of philosophical thought, to say that a mental state has propositional content is just to say that it represents the world as being a certain way. A proposition specifies the way the world would have to be in order for the state in question to be *satisfied*.[12] So, for

---

[12] See e.g., Searle, *Intentionality*. Searle has consistently defended the relevance of a larger Network of Intentional attitudes and a Background of non-Intentional capacities for fixing the propositional content of any single sentence

instance, what's required for the belief that there's beer in the fridge to be true is just that there now be beer in the fridge; the belief's propositional content is thus that there's beer in the fridge. An infinite range of otherwise distinct possible worlds are all compatible with that belief's being true, just so long as in each of those worlds there is beer in the fridge at the relevant moment. A proposition is also supposed to be the abstract 'thing' that's shared between two people (or the same person at different times) who say or believe (or take up other attitudes toward) 'the same thing' — whether they speak the same or different languages, whether they use the same or different words, whether they speak at the same or different times. Finally, on most views, we can determine by reflection — perhaps on the proposition's own internal structure, perhaps in some other way — at least some of the other propositions that must, could, or could not be satisfied in the same situations as the initial proposition. These modal relationships between propositions then generate rational, inferential connections between the mental states which have them as their contents. So far, all these claims are supposed to be ontologically neutral about just what propositions are, or about what the 'representing relation' consists in or requires.

One natural understanding of the idea that a mental state *lacks* propositional content, then, is that it does *not* represent the world as being any particular way, and that it stands in no rational, inferential relationships with other mental states. On this understanding, non-propositional mental states are modeled after, if not strictly equivalent to, feelings like twinges and tickles. A twinge isn't *about* anything else; it doesn't represent the world as being any way at all. Ethical expressivism à la Ayer and Stevenson treats ethical statements as non-propositional in this sense. Such statements specify no propositional content apart from the basic specification of the state of affairs being evaluated (as in 'It was wrong *that* she stole your ice cream'). Instead, they express a feeling the speaker has when confronted by certain sorts of situations, and they attempt to arouse the same sort of feeling in the hearer. Ethical statements mark the presence of certain causal correlations between encountering a situation and feeling a certain way, but they do not represent that situation as being any particular way.

---

or thought. *Relative* to the Network and Background, however, a proposition specifies just some conditions which must obtain in order for the state to be satisfied.

We clearly cannot accept these conclusions about what it means to be non-propositional if we want to make sense of 'seeing-as' as a non-propositional state while also retaining its representational character. On the non-representational reading of non-propositionality, seeing things under a certain aspect would be a distinctive sort of experience, one which happened to be aroused in certain sorts of situations, but not one that represented the *things seen* in that situation in any particular way. Such an analysis of seeing-as would exclude precisely what is so interesting about the actual perceptual phenomenon, and in turn what is so inviting about talk of 'seeing-as' and 'perspectives' in relation to metaphor: the idea that seeing-as is a way of experiencing and making sense of *the world*, and particular things in it.

Neither can we conclude on this basis, though, that seeing-as and aspectual thought must therefore be propositional. To accept this would also undermine their explanatory value. The whole point of talking about a 'perspective' in the first place was that a perspective cannot be reduced to simply entertaining a thought, in the standard there's-beer-in-the-fridge sense of that phrase. If a perspective or an aspect could be cashed out in propositional terms, then we should be able to state that propositional content directly, albeit perhaps laboriously. But it is precisely the intuition that this is *not* possible which drives people to talk about 'aspects' and 'perspectives' in this context.

We therefore need a different way of drawing the **propositional—non-propositional distinction**. A narrower understanding of propositions treats them as the sort of thing that sentences express, and as having a form and structure analogous to that of sentences.[13] The relevant contrast on this understanding, then, is between sentence-like ways of representing things and non-sentential, *pictorial* modes of representation. This seems to be the thought that Davidson is expressing when he writes,

> When we try to say what a metaphor 'means', we soon realize there is no end to what we want to mention...How many facts or propositions are conveyed by a photograph? None, an infinity, or one great unstatable fact? Bad question. A picture is not worth a thousand words, or any other number. Words are the wrong currency to exchange for a picture.[14]

---

[13] See e.g. Jeffery C. King, "Structured Propositions and Complex Predicates."
[14] "What Metaphors Mean," p. 263. In this chapter I leave aside Davidson's negative claim that *neither* pictures nor metaphors stand for or represent things as being a certain way, and instead try to make positive sense of the claim that metaphor represent *in a similar way* as pictures do.

Several theorists — most notably, Nelson Goodman — have attempted to articulate the fundamental differences between **sentential** and **pictorial** forms of representation. I will just mention the most salient differences here, and indicate briefly their potential relevance to aspectual thought.

The first major difference is that pictures, but not sentences, are *analog* modes of representation. This term itself has many definitions, but we can isolate two important qualities, **density** and **repleteness**, each of which can have both a syntactic and a semantic implementation. When a symbol system is *syntactically* dense, there is continuous variation in what counts as a different character: the smallest difference in the mark makes a difference to its symbolic significance. When a symbol system is *semantically* dense, there is continuous variation in what those marks can refer to or be about.[15] Density concerns continuity and hence representational *precision*. Many representational systems, such as analog watches and echocardiograms, are syntactically and semantically dense without being pictorial.

Repleteness, by contrast, concerns *richness* of representation; representational systems can be more or less replete. *Syntactic* repleteness requires that a relatively wide range of the mark's properties affect what it represents.[16] This distinguishes pictures from dense but uni-dimensional modes of representation such as watches and graphs. *Semantic* repleteness requires that the symbol represent a wide range of properties simultaneously: color, size, shape, location, texture, and so on. Paradigmatic forms of pictorial representation, like perception itself, exhibit all four features — both semantic and syntactic density and repleteness — to a high degree. Linguistic representation ranks low on all four counts; and many other forms of representation fall somewhere between the two extremes.

People have felt that the sort of understanding associated with metaphors ranks high on the scales of something very like both semantic density and semantic repleteness. Thus, with respect to density, they have often felt that metaphors communicate much more precise and vivid information than many literal statements do. We often turn to metaphor when we want to capture *just* how something looked or sounded or felt, when the usual coarse semantic categories won't do. So, for instance,

---

[15] See Nelson Goodman, *Languages of Art*, ch. IV, §§2,5.
[16] *Languages of Art*, ch. VI, §1.

(5) In the voice of a wheezing, tired old bagpipe, he began to recite his tale

conveys a more specific sense for the relevant sound than does

(5') In a thin, nasal tone of voice, he began to tell his story.

I will try, in the coming chapters, to do justice to what is right about the intuition that the understanding metaphors induce is often more precise than that of literal utterances. The problem with putting the point in this way (which is how it is often put) is that our language *does* have semantic resources for directly and literally specifying values in highly precise ways. Demonstrative reference is custom-tailored to specify semantic values, be they objects or properties, in as fine-grained a manner as our powers of perceptual discrimination allow. Further, we can and sometimes do introduce other forms of precise specification into our language on an *ad hoc* basis: numerical scales, for instance. Many of these forms of representation are not strictly speaking continuous, but they can still be employed to make semantic distinctions that are as fine-grained as we like. Thus, what we need to explain is not how metaphor conveys content that's more precise than ordinary semantic content, but how metaphor manages to accomplish something like demonstrative specification.[17]

The intuition that the sort of understanding engendered by metaphors is semantically *replete* is significantly stronger than the intuition that it is dense or precise; indeed, it seems to be one of the strongest motivations for talk of 'perspectives' and 'non-propositionality' in this context. As Davidson says, many metaphors don't just bring one thing to our attention. They invite us to keep on noticing more new features, of various kinds, as we spend more time with them. Here, we might take **Kant's definition of an *aesthetic idea*** as an apt description of the analogue in aspectual thought to the repleteness, and hence the openendedness, of perception and pictorial representation:

> By an aesthetic idea I mean that representation of the imagination which induces much thought, yet without the possibility of any definite thought whatever, i.e. *concept*, being adequate to it, and which language, consequently, can never get quite on level terms with or render completely intelligible.... In a word, the aesthetic idea is a representation of the imagination, annexed to a given concept, with which, in the free employment of the imagination, such a multiplicity of partial representations are bound up, that no

---

[17] I find this intuition to be the most compelling support for Stern's treatment of metaphorical meaning as analogous to demonstrative reference, as opposed to the purported evidence about syntactic constraints on metaphorical interpretation I discussed in chapter 1. See in particular his discussion of exemplification in chapter 5 of *Metaphor in Context*. I return to demonstrative concept formation in §3.5.1.

expression indicating a definite concept can be found for it — one which on that account allows a concept to be supplemented in thought by much that is indefinable in words.[18]

A key challenge of the coming chapters will be to do justice to this intuition in an explicit and relatively precise way. Further, as the Kantian definition reminds us, the phenomenon of aspectual thought is not limited to comprehending metaphors. At the very least, a rich engagement with fiction often involves cultivating a very similar sort of multifaceted and open-ended perspective. And in the right contexts, literal, true utterances can work similar effects. I want to explain what pattern of thought is common to all these uses of language — metaphor, fiction, and certain literal, true utterances — which makes Kant's description, and talk of non-propositional repleteness, apt. My account of aspectual thought is unlikely to capture all that people have wanted to mean when they have cited Kant's 'aesthetic ideas' and the 'free play of the imagination' — no one account could. And in any case, my primary task here is making sense of how metaphors work. However, I think that by focusing on a pattern of thought which is characteristic of metaphor but which is also exploited by other uses of language, we can gain a clearer understanding of what that pattern of thought is, of how metaphor relates to those other uses of language, and of why metaphor should be such a natural and pervasive form of communication. Armed with an account of this pattern of thought, we can then go on to analyze what is distinctive about the way that metaphor works.

There are a few further differences between pictorial and sentential forms of representation which bear mentioning here. The first is that pictures present their information in a non-compositional fashion, whereas sentences are semantically articulate.[19] That is, there is no analogue to the subject-predicate distinction in pictorial representation. A second point, corollary to this, is that there can be "no bare depiction of particulars," as Robert Hopkins says.[20] We have reason to think that singular referring terms, such as names or demonstratives, typically function like variables: they simply fix *which* object is their

---

[18] *The Critique of Judgement*, Part I, Book II, §49, pp. 175-179; emphasis in original.
[19] Cf. Moran, "Seeing and Believing," pp. 92-3.
[20] *Picture, Image, and Experience: A Philosophical Inquiry*, p. 24.

semantic value, without specifying anything about *how* that object is.[21] By contrast, in pictorial representation *what* is represented and *how* it is represented as being are bound up in a single unity. Because of this, and because pictorial representation is typically semantically dense, representing some thing at all necessarily involves representing it as being a particular, relatively determinate way.[22] The third point is that all pictorial representation depicts its content from a certain point of view: from down and to the left, for instance. I have not (yet) found any particularly deep or interesting analogies between these three facts about pictorial representation and our intuitions about aspectual thought, although some echoes of the second two points will crop up in chapter 4.

2.3: *Seeing-As*

So far, I've uncovered a few analogies between perception and the pattern of thinking which metaphors engender. In particular, I've noted that both phenomena make possible an open-ended, rich collection of thoughts, but don't seem to be reducible to those thoughts. But the original intuition I set out to explain was not just that metaphorical understanding is like perception or pictorial representation in general, but that it is like the phenomenon of seeing-as specifically. I now turn to that phenomenon. In the next two chapters, we'll find that nearly every point made here about seeing-as has an analogue in aspectual thought.

2.3.1: *Shifts of Aspect*

Two very familiar examples of aspect-perception are occasioned by the Gestalt duck/rabbit figure made famous by Wittgenstein, and the classic illusion depicting both a young lass and an old crone (Figures 1 and 2). What is so distinctive about these figures is that two dramatically different aspects are 'hidden' in a single figure. These aspects are mutually exclusive: if you're lucky, you can successfully

---

[21] Cf. Saul Kripke, *Naming and Necessity*. Titles, nicknames, and some complex demonstratives seem to be exceptions here.

[22] One might object here that sentences are really the minimal semantic unit, and that sentences do always specify both what is being talked about and how it being said to be. I am sympathetic to this position. The point remains, however, that these two facets are bound together more intimately in pictorial than in sentential representation.

FIGURES 1 AND 2 HERE

see either figure in either way, but you can never see either in both ways at once. Changing from one way of seeing to the other involves a 'gestalt shift'. The figures thus serve as particularly forceful examples of aspect-perception more generally, which need not always involve this stark incompatibility of aspects.

The relevant question for us is, what makes the difference between the two ways of seeing each figure? One might be inclined to think that because I'm continuously looking at a single figure, because the perceptual stimulus itself has not changed, therefore the difference must lie in my thoughts about or interpretation of that stimulus. But this is a highly **misleading** way to describe the phenomenon. Clearly, thought and interpretation have *something* to do with it. A change of aspect is typically accompanied by a change in which concepts seem appropriately to characterize one's perception, and in which descriptions one is inclined to offer of what one sees. But the difference between the two ways of seeing each figure isn't merely, or even primarily, one of interpretation. The crucial difference lies in the role that a thought or concept plays *in* the perceptual experience itself, not just in the thoughts we have *about* that experience or the objects experienced. As Wittgenstein says, seeing-as is not just "looking plus thinking."[23]

To grasp the difference between seeing-as and mere 'looking plus thinking', consider the old crone/young lass figure. Suppose Mary has only ever seen the figure as of an old woman. She knows *that* a picture of a young lass is embedded in the figure, and she thinks of the illustration in those terms (e.g., by thinking "There's that old crone/young lass figure I can never get"). But she still doesn't actually see it *as* a picture of a young woman. Someone may have pointed out for her what each element of the figure is supposed to represent (a nose here, a necklace there). Let us even suppose that she accepts on reliable testimonial grounds all the propositions true of the figure. Nonetheless, having those thoughts about the figure while also looking at it is not itself sufficient for seeing it that way — as anyone who has struggled to see the girl in the picture realizes.

When Mary finally does undergo a gestalt shift and succeed in seeing it as a picture of a young lass, the difference between her two experiences does not lie in any new proposition she has come to believe (by hypothesis, she believed them all already), nor in her acquiring or bringing to mind any new

---

[23] *Philosophical Investigations,* II.xi, p. 211.

concept.[24] Even so, there is a real difference between her old and new experiences. As Wittgenstein says, "I *see* that it has not changed; and yet I see it differently"[25]: a new *aspect* has dawned in the perceiver's experience.

This difference is most aptly described as a **structural** one. The figure is now organized for Mary in a new way, so that "parts of the picture go together which before did not."[26] (For instance, the parts which serve as the old crone's mouth and nose no longer 'go together' when the picture is seen under the aspect of the young lass; they have become the lass's chin and necklace.) This new organization in turn assigns new *roles* to elements within the picture. As a result, it also acts as a 'filter', *highlighting* those features that play crucial roles and downplaying those that do not. (For instance, the relatively inconspicuous wart on the old woman's nose now plays a crucial and prominent role in the young woman's face. Smoothing out that line would likely undermine our ability to see the young lass in the figure; it would not markedly affect our overall sense of the old woman's face.) And having this new organization available doesn't just alter our current experience; it also changes the way we *go on* to engage with the figure. Instead of "reading it like a blueprint" — identifying each constituent feature separately and recalling what it was supposed to stand for — we now "know our way about"[27] the figure as a whole. That is, the figure now makes overall sense to us, and we intuitively think of and deal with constituent features in terms of their role within the whole.

So far, I've been arguing that the dawning of an aspect in perception involves a structural, and not primarily a propositional, change. This structural change imposes new relations between elements, highlights some and not others, and alters our future engagement with the figure. However, perhaps this insistence on a structural as opposed to a propositional difference is too strong. Perhaps we can say that even though Mary acquires no new proposition, still the aspect still involves a sort of propositional change after all: a change of *attitude* toward a proposition. Before the aspect dawned, Mary merely

---

[24] In particular, by hypothesis she already believed and had in mind all the relevant propositions of the form '*This* is her nose', '*That* is her necklace,' etc.
[25] *Philosophical Investigations,* II.xi, p. 193.
[26] *Philosophical Investigations,* II.xi, p. 208.
[27] *Philosophical Investigations,* II.xi, p. 204.

believed that the figure could be seen as of a young woman, or that this particular feature represented the young lass's nose. By contrast, one might argue, she now *sees that* these same propositions are true. She now has a perceptual, and not just a testimonial, source of evidence for her beliefs; and this gives her a new reason for believing them and other, related propositions.

This point is fine, as far as it goes; the difficulty is that it does not go far enough to undermine the point about there being a structural difference. First, the difference between testimony and perception is not just one of checking off a different box on an 'evidential basis' form in our bank of beliefs. There is something about perceptual experience which makes it an especially compelling source of evidence. I suspect (though I can't argue here) that we will not be able to explain why this is without appealing to *how* we see what we do — to the fact that the concepts receive their instantiation in such a particular way before our very eyes.[28] More specifically to the point, however, merely saying that Mary now *sees that* the figure is a young-woman picture leaves mysterious what it is about her experience that has changed. Typical cases of coming to see that *P* involve noticing some previously overlooked fact: that there are fingerprints on the glass, for instance. But by hypothesis, in this case Mary notices no new fact. Rather, she comes to see the same old facts in a new *way*. And in order to explain what *this* involves, we still need to appeal to the changed organizational structure of her perceptual experience, and not just to her different mode of access to a proposition already believed.

2.3.2: *All Seeing is Seeing-As*

I've focused so far on a quite special case of perception: the sudden dawning of a new aspect for seeing an intentionally ambiguous figure. Similar experiences also occur, of course, in less craftily ambiguous situations, and in relation to scenes which are not pictures. Thus, I may struggle to hear the return of the theme, now inverted, in a symphony; to see the anguish and anger that I'm told is expressed

---

[28] I think this at least points to an important lacuna in the causation-of-belief analyses of perception advocated by Robert Brandom in *Making it Explicit*, and by Davidson in "A Coherence Theory of Truth and Knowledge." I also think attention to seeing-as is what is required to render McDowell's talk in *Mind and World* of perception as 'wringing commitments' from us less mysterious. See my "The Conceptual but Non-Propositional Content of Perception."

by a De Kooning abstract painting; to see my all-too-brave bereaved friend's face and posture as marked

by grief; or to see the molecular structure of a chemical solution under the microscope. In all these cases,

my experience itself is altered if I do succeed in applying the concept within the perception.

Wittgenstein himself maintained that talk of 'seeing-as' was appropriate only to cases in which

there is a relevant potential *shift* of aspect. As he says,

> One doesn't 'take' what one knows as the cutlery at a meal *for* cutlery; any more than one ordinarily tries
> to move one's mouth as one eats, or aims at moving it. If you say "Now it's a face for me", we can ask:
> "What change are you alluding to?"[29]

As a point about conversational felicity, Wittgenstein's comment is surely apt. But he seems to have

taken it to show that ordinary seeing is *just* seeing *simpliciter*. Others have wanted to argue against this

that most or all of the seeing that we normal, engaged, adult humans do is seeing-as. That is, thinking

about dramatic shifts of aspect, whether or not the scene in question has been intentionally designed to

encourage such shifts, helps us to notice that we were already seeing the scene *in a certain way* before the

shift occurred, and invites us to generalize this point to cases in which no such shift does in fact occur.[30]

Paul Churchland brings out the "plasticity of perception" and so ultimately of our conceptions of

reality, even in those arenas where we never before considered a shift of aspect, with the following real-

life thought experiment. First he imagines a society of people who operate with modern physical theory

as their everyday conceptual scheme:

> These people do not sit on the beach and listen to the steady roar of the pounding surf. They sit on the
> beach and listen to the aperiodic atmospheric compression waves produced as the coherent energy of the
> ocean waves is audibly redistributed in the chaotic turbulence of the shallows...For these people the world
> is seen (felt, heard, etc.) to cohere, even in its prosaic details, in ways to which our ordinary conceptions are
> utterly blind.

Churchland recognizes that these people may seem, at least initially, quite foreign to us, and so that his

description of their perceptual experience may not be compelling. In order to bring the possibility of such

an 'embodied' everyday scientific understanding to life, he suggests that we ourselves undertake to *see*

the nighttime sky *as* Copernican theory describes it. Up to this point, he says, "Our minds, perhaps, have

been freed from the tyranny of a flat immobile Earth, but our *eyes* remain in bondage." But if we invest

---

[29] *Philosophical Investigations,* II.xi, p. 195.
[30] See e.g., Searle, *Intentionality*, ch. 2.

the effort required to follow the instructions he lays out, then we can bring the much more sophisticated and informative set of concepts articulated in Copernican theory to bear *within* our perceptions. Churchland concludes: "I urge the reader not to judge the matter from my own spare sketches. Judge it in the flesh some suitably planeted twilight. A vertiginous feeling will signal success."[31] C. S. Lewis makes much the same point about the "vertiginous" effect of bringing the medieval conception of the universe to bear in our perception of the night sky.[32] But if we can come to recognize that we have been seeing something as pervasive and familiar as the night sky *in a certain way* all along, and if we can experience the shift to a new way of seeing it, then surely seeing-as must be a quite general phenomenon.

As these examples bring out, for normal, engaged, human adults it is typically the concepts and beliefs we deploy that make the crucial difference between different perceptual experiences of the same thing. We normally treat our perceptions as perceptions *of* something, and *what* we take our perceptions to be of partly determines *how* we see it: <mark>what we see it *as*.</mark> Even in situations where we 'bracket' our usual object-oriented concepts, such as when we focus on the colors and textures of what we see, or in which we don't know just what it is we're seeing, we still treat our perceptions as *of*, say, a bright-red silky-smooth something, or *of* a strange lumbering animal.

Indeed, it seems that this point can be generalized. Many philosophers have thought for Kantian reasons that the very intentionality of perception itself — its being object-directed as opposed to being a mere passing sequence of sensations — requires that we bring some concepts to bear in, or at least on, perception. As Strawson says,

> ...there would be no question of counting any transient perception as a perception of an enduring and distinct object unless we were prepared or ready to count some different perceptions as perceptions of one and the same enduring and distinct object. The thought of other actual or possible perceptions as related in

---

[31] *Scientific Realism and the Plasticity of Mind*, pp. 29-34.

[32] "These facts [about the distances between planetary bodies assumed under a medieval cosmology] are in themselves curiosities of mediocre interest. They become valuable only insofar as they enable us to enter more fully into the consciousness of our ancestors by realizing how such a universe must have affected those who believed in it. The recipe for such realization is not the study of books. You must go out on a starry night and walk about for half an hour trying to see the sky in terms of the old cosmology. Remember that you now have an absolute Up and Down. The Earth is really the center, really the lowest place; movement to it from whatever direction is downward movement. As a modern, you located the stars at a great distance. For distance you must now substitute that very special, and far less abstract, sort of distance which we call height; height, which speaks immediately to our muscles and nerves. The Medieval Model is vertiginous...." (*The Discarded Image*, p. 98).

this way to the present perception has thus a peculiarly intimate relation to our counting or taking — to our ability to count or take — this present perception as the perception of such an object.[33]

In order really to be perceiving an object-filled world, the thought goes, the subject must be prepared to categorize her perception at this moment within a structure of other perceptions, both of the same object and of qualitatively similar, though numerically distinct, objects. Only then does she represent the world as containing *objects*: that is, persisting entities with multiple properties which can change through time, and with which she interacts on different occasions. This is tantamount to saying that the subject must at least deploy the concept *object* relative to her perceptual experiences. Of course, ordinarily we don't just take our perceptions to be recurrences of this or that object *in general*, but to be of a certain *sort* of object: a tree, or car, or whatever. But this just means that we deploy finer-grained sortal concepts to classify the sorts of objects we take our perceptions to be of. If this Kantian line of thought is correct, then there is no basic, pure, non-conceptual way in which we could, let alone normally do, see the world around us.

There are many important and interesting questions lurking here about how to understand the requirements of concept-possession and the content of other beings' perceptual experiences. In particular, this Kantian view makes it difficult to account for the perceptual lives of non-human animals and infants. Many philosophers argue that such beings do not have conceptual capacities strictly speaking, on the grounds that mere categorization, even systematic categorization of the sort Strawson describes above, does not yet count as genuine conceptualization.[34] In addition, they argue, the subject must be able to comprehend the possibility that this particular way of categorizing what she sees might be in error. That is, she must be able to make sense of the possibility that what she takes to be a tree might not in fact be a tree. Perhaps she must even make sense of the possibility that what she takes to be a perceptual experience of an object might in fact be a mere hallucinatory sensation. This stronger requirement on genuine concept-possession is compelling, because it respects the intimate connection that seems to obtain between concepts, judgment, and epistemic responsibility. The disadvantage of it is that it seems

---

[33] "Imagination and Perception," p. 52.
[34] Cf. Wilfred Sellars, *Empiricism and the Philosophy of Mind*; McDowell, *Mind and World*.

to consign non-human animals and infant humans to experiencing mere sensations rather than perceptions of the world, because they seem not to meet the requirement.

I myself am inclined to think that we need a more nuanced, graded way of thinking about concepts, or at least about 'proto-concepts', if we are to make sense of the mental lives of animals and infants. On the one hand, we should respect the important differences between our conceptual capacities and the representational faculties of relatively sophisticated animals like primates, parrots, and household pets. But on the other hand, we should not simply lump these animals together with algae and thermostats as mere differential responders, either.[35] Such animals do seem to track objects in the world, to be capable of recognizing *the same thing* over again, to learn new behavior in the light of repeated past experiences with both particular individuals and general types of entities.[36] Further, it seems that at least some animals do see things in a certain *way*, in that they can make some distinction between two extensionally equivalent but intensionally distinct modes of presentation. For instance, dogs and pigeons seem to be capable of being trained to respond differently to squares and diamonds.[37]

However, these further questions are irrelevant for my current purposes. Again, my only concern here is tracing out those features of perceptual experience which provide relevant analogies for understanding the comprehension of metaphors. Since metaphorical comprehension clearly requires possessing concepts in a full-blown sense of that term, I need only be concerned with the representational capacities of concept-wielding subjects. For these subjects, it does seem to be true that concepts (where this includes demonstrative concepts like 'that shape') play an essential role in normal perceptual experience. There may also be features of perceptual experience which are essentially non-conceptual: *qualia*, for instance. I need not take a position on this matter here, because there is no analogue to *qualia* in aspectual thought.

---

[35] Cf. Fred Dretske, "The Nature of Thought"; Ruth Garrett Millikan, "Thoughts without Laws."
[36] Cf. Marc Hauser, *Wild Minds*.
[37] Cf. Christopher Peacock, "Scenarios, Concepts, and Perception."

2.3.3: *The Two Axes*

For the claim that all seeing is seeing-as to be phenomenologically plausible, it's important not to assume that all seeing must be conceptual in just the same way. In particular, we must insist that conceptuality does not require the self-conscious, active application of concepts that we find with the illusion figures. Usually, we simply see, and make sense of what we see, without much fuss. Rather than dividing into two distinct classes of seeing *simpliciter* and seeing-as, as Wittgenstein seems to have thought, perceptual experience varies along a continuum, according to the ease and immediacy with which we bring the relevant concepts to bear in perception.

Thus, for instance, I apply the concepts *knife* and *fork* quite automatically when I encounter my dinner cutlery each evening. The relevant concepts may even be active in anticipation of the relevant perception, as when I hear my friend's glamorous new Lexus coming round the corner and so am predisposed to see the car as a Lexus. In addition to these quite immediate applications, Richard Wollheim distinguishes at least three further stages of increasingly delayed and effortful application. First, I may be looking at a scene without entertaining the relevant concept, but when that concept is suggested to me by someone else, it immediately becomes active in the perception: I respond with something like, "Ah yes, now of course I see *a* as an *F*." Next, the concept may join in the perception only "after long scrutiny." Here, I scrutinize the object in order to figure out whether it is an *F*, or to figure out what sort of thing it is within a broader class of which *F* is one sort. For instance, I might want to figure out whether this tree is an oak, or just what sort of tree it is. In the end, though, the oakiness of the tree simply 'clicks into place.' Lastly, I may be able to join the concept to the perception only through "an act of will on my part."[38] In this case, I may want, say, to see the tree as an oak, but manage to do so only by actively imposing the concept on the perception.

Attending to the whole range of cases along this continuum brings out the odd sense in which we are both active and passive in perception.[39] As we move out along the axis, 'seeing-as' becomes an

---

[38] "Seeing-as, Seeing-in, and Pictorial Representation," p. 221.
[39] Cf. McDowell, *Mind and World*.

increasingly reflective process, increasingly a matter of our voluntary control. At the immediate end of

the scale, even when I see a figure or scene quite automatically in a certain way, it still makes sense for

me to *try* to see it in a different way, and for someone to request that I do so.[40] Any new bit of

information about which feature in the figure or scene is supposed to correspond to which element of the

aspect may help me to comply with such a request. But at the same time, there is no guarantee that any

such information, or any way of training my eyes, or anything else, will in fact induce the relevant aspect.

That is, the aspect's dawning is never entirely within my control. I can't, for instance, sensibly be ordered

to see a scene in a certain way; rather, I must allow the aspect to dawn on me. This fact becomes

dramatically apparent with more complex optical illusions.[41] In these cases it may take quite a long time

to see the scene as desired, and I may never succeed. When I do succeed, I am rarely sure just what led to

my success.

In addition to what we might call the axis of *immediacy*, Wollheim identifies another axis along

which perceptual experiences vary. This axis concerns the sort of *commitment* we have to the

applicability of the concepts we bring to bear. In the paradigmatic case of engaged perception, belief in

the concept's applicability follows directly upon, or is even simultaneous with, perception: in these cases,

seeing *is* believing (although as we have seen, it is not just that). Further out along the axis come "likely

supposition, informed guess, outside bet," as for instance, when I guess, with more or less certainty, that

the animal out in the field is an antelope. Finally, as Wollheim says, there is "the case where there is no

commitment at all to the satisfaction of the concept by the object, and imagination or make-believe takes

over."[42] A paradigm case of this might be seeing a mountain as the body of a sleeping giant.

It is important to note here that even when I have no inclination at all to believe that a given

concept truly applies to the perceived object, still there are genuine constraints on which concepts I can

bring to bear in my perception of it, and so on how I can see the object in question. Even the most

---

[40] Cf. Roger Scruton, *Art and Imagination*, p. 95.
[41] See http://www.geocities.com/SouthBeach/Port/2701/illusions.html for a relatively comprehensive collection of
optical illusions, including some particularly daunting ones.
[42] "Seeing-as, Seeing-in, and Pictorial Representation," p. 221.

monumental and patient effort will not enable me to see a placid backyard pond as a snarling, snapping tiger. So even within the realm of perceptual make-believe, not just anything goes. Indeed, even at the extreme outer limit of the continuum, there is still a sort of commitment to the concept's applicability to the object. Although there is no commitment to the truth of the concept's applying, I must still be committed to the concept's being *appropriate* for that object in some way. One task for the next chapter is to make sense of what this non-truth-directed appropriateness could be, both for aspectual thought and for perception.

Often, perceptions vary in a parallel fashion along the two axes. That is, perceptions in which concepts find easy, immediate application tend to be perceptions for which I fully accept that the concept actually applies. And perceptions for which only an act of will allows me to render the concept active tend to be perceptions in which I am merely trying out or imagining *a* to be *F*. However, a perception's location on one axis does not always correlate with its position on the other. Thus, for instance, I may easily and immediately see a cloud as a camel, or a tree stump as an old man with a walking stick, even though I am under no illusion that either one really is, or is even a representation of, either of those things. Sometimes such concept-applications can be hard to shake, even when we want to rid ourselves of them. So, for instance, I might continue to see my colleague as resembling a particularly hated figure from my past despite my best efforts to perceive him differently. At the other extreme of 'mismatch' along the axes, a scientist might believe for independent, conclusive reasons that a smear on a microscope slide is a sample of a certain cell type, and desperately want to verify that it is so that he can publish the study he's been working on for many years. He may still only be able to see the characteristic pattern on the slide after considerable effort, if at all.

2.4: *What Seeing-As Misses*

The wide range of cases we've been considering so far do provide good evidence, I think, that concepts normally play a role *in* perception and not just in our thoughts *about* our perceptions. Thus, we have good reason to accept that nearly all of our visual perception (again, where this means the perception

of normal, engaged, adult humans) involves seeing-as.[43] However, I have been glossing over two further complications, whose analogues will have crucial implications for aspectual thought in the coming chapters.

2.4.1: *Not all Seeing is Seeing Under an Aspect*

When people talk about the sort of phenomenon I've been discussing, they usually employ the phrases 'seeing-as', 'seeing x in terms of y', 'seeing x in a certain way', 'seeing x under an aspect', and 'seeing x in a certain light' roughly interchangeably. And they often explain the phenomenon, as I have here, as one in which a concept is brought to bear on and in the perception of an object. But this easy equivalence of phrases, and this general description, are not fully accurate — at least not as I will employ the terms. On my use of the terms 'seeing-as' and 'seeing under and aspect', one can see a as F without there being an aspect under which one sees a; so aspects are not necessary for seeing-as. And conversely, one can see a under a certain aspect without that aspect being aptly characterizable in terms of a concept F that one brings to bear in perception; so seeing-as is not always sufficient to capture an aspect.

On the first point: for any perception, I will be prepared to predicate a host of concepts of what I see; I see those things *as* instantiations of those concepts. Further, my perception is differentiated from that of other people partly in virtue of the different concepts we each possess and deploy. But none of this requires the holistic organization that is so distinctive of aspects. Thus, for instance, when I see the old crone/young lass figure as a picture of a young lass, I do see her *as* wearing a hat, but my perception is not under the *hat* aspect: that concept informs just one element within the overall aspect. The concept *young lass*, by contrast, does impose an overall structure to the figure, so that it 'hangs together' for me in a new way. Or, to take a non-pictorial case, I may see the scene before me under the overall aspect of domestic harmony, while seeing just the girl doing her homework as an instantiation of the concept *girl*

---

[43] Visual perception is the richest and most heavily concept-laden of our sense modalities. Hearing seems to run a close second; it is less clear that the other modalities are always conceptualized. I think more attention should be paid to the varieties of experience across the modalities than philosophers typically do pay, but I've fallen into the usual narrowness here.

doing her homework. As Wollheim says, seeing-as always manifests some degree of *localization*;[44] while aspectual seeing always encompasses multiple elements within a field. Thus, 'seeing-as' has weaker criteria of attribution, and so a broader range of application, than 'seeing under an aspect' does.

When people invoke 'seeing-as' in relation to metaphor, it is the latter phenomenon that they are really after, I think. They mean to capture the sense in which a metaphor can reorganize our overall understanding of an individual or situation. When I talk about 'aspects' and 'aspectual thought' in the coming chapters, I will always be using 'aspect' to mean a way of thinking or perceiving which imposes this sort of overall holistic organization. It will be absolutely crucial to keep this definition in mind: as I employ the term 'aspect', aspects impose an overall holistic organization, while mere seeing-as or thinking-as need not do so.[45] The task of the coming chapters will be to make sense of how aspects, applied in thought, manage to accomplish this.

The difference between aspects and seeing-as becomes more apparent if we attend to the fact that some aspects can organize a perceptual whole without any single feature within the perception being seen *as* instantiating them at all. In these cases, localization fails entirely. So, for example, a scene may look threatening without any element in it being identifiable as the threat. Moods provide the most striking example of this: when I feel joyful, everything around me seems to be brimming with energy, but I need see no one element as an instance of joy. One might think here also of the sense of an impending storm: the storm's presence is at once everywhere and nowhere in my environment, and a certain light, a certain sort of quiet or of wind, pervade the entire scene. Similarly, music can 'color' our overall perception of a movie scene without applying to any element within it.

This leads us to the second point: that not all aspects seem to be specifiable in terms of concepts. At least on a *prima facie* understanding, moods and musical coloring crucially involve, not the application

---

[44] "Seeing-as, Seeing-in, and Pictorial Representation," p. 211.

[45] I recognize that drawing a distinction in this way between seeing-as and seeing-under-an-aspect does not fully accord with our ordinary linguistic usage. In particular, we sometimes use 'aspect' simply to talk about the angle from which we perceive something, or the facet of it that we perceive. This is also the sort of use which corresponds to the Fregean notion of a 'mode of presentation'. My use of 'aspect', here and in the coming chapters, is a semi-technical one designed to meet my own theoretical needs. I do maintain, though, that it represents a natural extension from our ordinary use.

of concepts, but the projection of emotions.[46] Thus, even putting aside more metaphysical concerns about

whether, say, evaluative terms express genuine concepts and so in turn about whether evaluative aspects

involve applying concepts,[47] still we cannot treat all aspects in terms of concept-application.

There is also another way, apart from moods and music, in which an aspect may crucially involve

applying something other than a concept, at least as we often think of concepts. Some aspects seem to

consist *entirely* in organizing what is perceived in a certain way. They attribute no property to the scene

seen or any element within it; nor do they take the scene as manifesting any emotion or attitude. As an

example of such an aspect, Roger Scruton offers (at least some experiences of) listening to music:

> Hearing a sequence as a melody is more like seeing a group of lines as a pattern...than it is like seeing a
> pattern of lines as a face. Here the 'organisation' of experience cannot be described in terms of the
> application of some independently specifiable concept.... the element of thought has been reduced to
> something entirely formal. There is no way of achieving even a partial description of the *content* of the
> musical thought: we can only point once more to the experience in which it is 'embodied'. And yet the
> formal properties of thought remain, removing the experience of music from the realm of merely animal
> mentality.[48]

When Scruton talks about 'seeing a group of lines as a pattern', he has in mind the different

experiences produced by seeing, say, a grid of multicolored dots as organized in rows, or as organized in

columns, or as organized diagonally. When we hear a sequence of notes as a melody, we do hear it in the

ways distinctive of aspectual perception. The notes 'hang together' with a certain contour; some notes are

more prominent than others; some notes lead up to and or fall away from the 'central' ones. Skipping or

changing some, 'ornamental' notes wouldn't change the melody significantly; doing the same to others

would destroy the melody entirely. But all of this seems to reflect just the organizational structure of the

perceptual experience itself. None of it attributes any further significance to the notes, as the application

of a concept would.

Some philosophers might argue that these 'purely formal' or structural aspects make no real

representational difference to how the scene is seen.[49] We still have all the same elements, arranged in all

the same genuine, objective relations, and it's not clear that there are any inferential consequences to

---

[46] Cf. Wollheim, "Correspondence, Projective Properties, and Expression."
[47] Cf. McDowell, "Values and Secondary Qualities."
[48] *Art and Imagination*, p. 176.
[49] John Searle, personal communication.

applying the one structure as opposed to the other. Others might argue that these aspects do involve

applying a concept after all: a purely formal or structural concept.[50] They might defend this by arguing

that there *are* inferential consequences to applying one such concept versus another. In particular, which

other objects or scenes count as *the same* in the relevant respect will be different in each case. Still other

philosophers might acknowledge that there is a real representational difference here, but insist that it is

not a genuinely conceptual one.[51] Again, although this issue is essential to an adequate account of the

role concepts play in perception *per se*, I do not think I need to settle it here. What matters is that we

notice that some aspects fail to fit the model of applying a standing, substantive concept like *old woman*

or *tree* or even *that color of red*, but do nonetheless still make a significant structural difference in our

perceptual experience. In §4.7.1, I return to these puzzling, purportedly 'non-conceptual' perceptual

aspects in light of my account of aspectual thought.

2.4.2: *Twofoldness and Seeing-In*

   I want to bring attention to one final distinction among perceptual experiences which tends to be

obscured by focusing on seeing-as in general, and with it one final feature of some perceptual experiences

which will be relevant for aspectual thought.[52] In my discussion of seeing-as and aspects thus far, I have

indiscriminately employed as examples both perceptions of representations and perceptions of the world

around us. Wollheim, by contrast, makes a sharp distinction between the two sorts of cases. The

perception appropriate to pictorial representations, he claims, must be analyzed in terms of seeing-*in*, not

(just) in terms of seeing-*as*. We see Napoleon *in* the painting, but we do not see the painting *as*

Napoleon. If seeing-as is relevant here at all, it consists in the fact that we see the painting *as* a painting

of Napoleon. But in order to explain this latter point, Wollheim argues, we need to appeal in turn to

seeing-in. We have already indirectly encountered one reason for making such a distinction between

---

[50] John MacFarlane, personal communication.
[51] Cf. Christopher Peacocke, "Scenarios, Concepts, and Perception"; "Does Perception Have a Non-Conceptual Content?"
[52] Thanks to David Davies for bringing this point back to my attention when I could actually make sense of it.

seeing-as and seeing-in: seeing-as requires localization, while seeing-in, like seeing under an aspect, does not. I can see a threatening storm, or that a villain has just departed, *in* a picture, without seeing any part *of* the picture *as* the threatening storm or *as* the just-occurred departure of the villain.

A second reason for making this distinction is that the sort of seeing appropriate to pictorial representation requires what Wollheim calls '*twofoldness*', which seeing-as precludes.[53] Twofoldness occurs when my visual attention is directed simultaneously to the medium of representation and to what is represented. If I am looking at a pictorial representation in the appropriate way, then my attention is fully devoted *neither* just to the painted surface itself (as it would be if I were studying the thickness of its brush strokes, or examining the surface of a photo for fingerprints), *nor* to the scene depicted (as it would be if I were using a pattern of clouds as an occasion to imagine a procession of menagerie animals). Rather, my attention is directed *to* the painted surface *as a medium* for seeing what is depicted.

This phenomenon captures another important component of what people are after when they invoke 'seeing-as' in talking about metaphor — although if Wollheim is right, they are not employing the appropriate analogy for capturing it. Metaphorical understanding crucially involves a sort of doubleness. To really 'get' a metaphor, it's not enough simply to recover the speaker's intended propositional content, or even to think of the object described in the desired way. One must also attend to how the conceptual scheme evoked by the literal meaning of the sentence uttered at once underwrites and is distinct from the way of thinking about the object which it produces. An adequate account of aspectual thought must make sense of what this 'doubleness' involves.

2.5: *Conclusion*

In this chapter, I've brought out a series of points about seeing-as and closely related perceptual phenomena. Three points deserve to be singled out for special attention. First, as I use the term, aspects involve the imposition of an overall holistic organization, not just the application of a concept to a particular element in the visual scene. In this sense, aspectual perception needs to be distinguished from

---

[53] "Seeing-as, Seeing-in, and Pictorial Representation," p. 213.

seeing-as. **Second**, aspectual perception is not just 'looking plus thinking': in aspectual perception, concepts (and sometimes other psychological constructs, such as emotions or simply formal patterns) play a structuring role within perception itself. As a result of this structure, the scene seen 'hangs together' in a certain way, and constituent elements take on a certain significance and certain mutual relations. **Third**, bringing a concept to bear in perception does not require believing that the object(s) seen actually instantiate the concept in question. But even when there is no commitment to the truth of the concept's applicability, still there are real limits on how we can see something.

I also identified a sense in which perception can be understood to be non-propositional without therefore being non-representational: perception (and pictorial representation) is more precise and richer than the thoughts typically expressed by sentences. Metaphors too often communicate a quite precise understanding; in this respect they can function like demonstrative terms.[54] The richness and openendedness of perception finds an analogous manifestation in the Kantian notion of an aesthetic idea which engenders the 'free play of the imagination'. Finally, I noted the 'twofoldness' that is exhibited by seeing-in, albeit not by seeing-as itself, and I suggested that this is one of the features in which people have been especially interested when they talk about metaphor in terms of 'seeing-as'.

With these resources for making sense of aspectual perception in hand, I turn now to developing a theory of aspectual thought.

---

[54] Recall that I said the claim that perception differed from propositional thought in being more precise was inaccurate because demonstrative terms specify contents with values as precise as our powers of discrimination allow.

CHAPTER 3: CHARACTERIZATIONS

3.1: *The Condition of the Very Possibility of Aspectual Thought*

In chapter 2, I surveyed the phenomenon of aspectual perception in order to understand why so many people have found it so natural to talk about metaphor in terms of 'seeing-as', 'perspectives', and 'seeing under an aspect'. The most important conclusion of that survey was that in aspectual perception, we don't just find 'looking plus thinking', but thought *in* perception. That is, in typical cases of aspectual perception, concepts and thoughts play a structuring role in perception by imposing an overall, holistic organization on our perceptual experience. This is itself, I think, an important insight. It's also an insight that's difficult for standard philosophical theories of mind to assimilate, because they typically lack the resources to make sense of concepts or thoughts playing a *structuring* role at all, as opposed to an inferential or dispositional one.

But my primary topic is not perception, but metaphor. And in chapter 2 we also saw that talk of 'seeing-as' and 'seeing under an aspect' applies only metaphorically to metaphor. In making sense of metaphorical comprehension, we are really interested in what I've called 'aspectual thought'. The most important point of commonality between aspectual thought and aspectual perception, I have said, is that thought plays a structuring role in each. In order to carry over the central lesson from aspectual perception to aspectual thought, then, we will need to make sense of 'thought *in* thought': of one thought restructuring another by imposing an overall holistic organization on it.[1]

However, theorists who investigate the sort of thinking associated with metaphor typically provide an account of 'thought *plus* thought' instead: of the juxtaposition of two thoughts. It's not surprising that they do this, because it's even more difficult for typical theories of mind to make sense of 'thought in thought' than of thought in perception. But as we'll see at several points in the next three

---

[1] Recall that I also cited three other important commonalities: both aspectual perception and aspectual thought are precise and rich; neither requires belief in the veridical applicability of the concept to the object seen or thought about; and aspectual thought, like the seeing-in that is appropriate to viewing pictorial representations, involves a sort of 'twofoldness'. The first two points will be addressed in this chapter, the third in the next.

chapters, an account that relies upon juxtaposition is inadequate to explain aspectual thought in general and metaphorical comprehension in particular.

Standard theories of mind are exclusively concerned with propositional attitudes and the relations between them. But propositional attitudes themselves don't have the right sort of structure to be *re*structured in the relevant way. That is, a propositional attitude is 'composed' (more or less literally, depending on the theory) of an attitude and a proposition toward which it is directed: for instance, a belief that *P*, or a desire that *Q*. But neither the attitude itself nor its relation to its propositional content seem to have any structure of the relevant sort.

The propositional content itself — or (depending on the theory) the mental representation of it — is in turn standardly taken to be 'composed' of concepts. Here we do find a complex structure, but again not one that can be restructured in the appropriate way. On most views, concepts themselves don't have any internal structure.[2] The structure of whole thoughts is taken to be fully determined by the mode of combination of their constituent concepts, on analogy with the syntactic structure of sentences. A new mode of combination for the same constituent concepts produces a totally new thought, typically about some other state of affairs altogether, on analogy with what happens when we restructure the sentence

(1) Confusion now hath made his masterpiece!

to produce

(2) Now his masterpiece hath made confusion!

But this isn't at all what we're seeking. When we think of Juliet as the sun, or of man as a wolf, we think about the same subject, and indeed often the same state of affairs (say, Juliet's being beautiful), in a new *way*. We don't think about something else entirely. An aspectual shift can produce *some* change in content, but there must be some substantive continuity as well.

In order to find the right sort of complex structure for aspectual thought, we need at least to ascend from individual propositional attitudes to structured relations among them. But here, standard

---

[2] They may be essentially associated with a structure — for instance, an inferential role. But making this structure internal would threaten compositionality. See Gil Harman, "(Nonsolipsistic) Conceptual Role Semantics."

theories of mind only recognize two sorts of relations that might link propositional attitudes. First, there are rational, inferential relations between attitudes in virtue of their propositional contents. These are exemplified, for instance, by explicit means—end reasoning about how best to satisfy one's desires. Second, there are non-rational, more or less idiosyncratic, merely causal associative relations to other attitudes and to possibly non-propositional states, like sensations and imagistic memories. These are exemplified by Proust's bite of madeleine bringing to mind thoughts and memories of his grandmother.

Because on many views the inferential relations between propositions at least partially determine *what* a given proposition represents, we can't restructure those relations without again thereby altering *what* is being thought about. Thus, these too are not viable candidates for explaining 'thought in thought'. The most plausible available candidates for explaining how aspectual thought can restructure our thoughts would thus seem to be the merely causal associations between attitudes. Perhaps when we think of something under a new aspect, we restructure which thoughts first *remind us* of which other thoughts and images. This is, I think, the sort of story Davidson has in mind.

Max Black gives us a lead on how this might work in the following passage:

> what is needed [for interpreting a metaphor] is not so much that the reader shall know the standard dictionary meaning... as that he shall know what I will call the *system of associated commonplaces*...From the expert's standpoint, the system of commonplaces may include half-truths or downright mistakes...the important thing for the metaphor's effectiveness is not that the commonplaces shall be true, but that they should be readily and freely evoked.[3]

Black's main point here is that the interpretation of 'Sam is a pig' or 'Bill is a gorilla' depends not so much on what we think is *true* about pigs and gorillas, but on something like our *stereotypes* of them. I will spend some time in §3.4 cashing out what I think is right, and important, about this point in particular. What is relevant right now is Black's suggestion that we seek a complex *system* or pattern of thoughts that are associated in some important way. I think this is on precisely the right track.

What Black doesn't give us, however, and what we need, are some substantive criteria for determining what a 'commonplace' is, and perhaps more importantly, how these commonplaces are 'associated' into a coherent system. At the very least, we will need to say a lot more about what these

---

[3] "Metaphor," p. 40.

relations are like, so that we have a clear sense of what the relevant sort of structure among thoughts might be. The primary task of this chapter is to do just that.

I will depart from Black in two important respects, however. First, I think that the pattern of thought which underlies metaphorical thinking — aspectual thought — is also employed more generally, though in somewhat different ways, in our thought and communication. I therefore want my account of the 'system of associated commonplaces' to work for more than just the comprehension of metaphor. Second, I will argue that once we get clear on what these relations are, we will also see that a merely causal associationist account is inadequate to explain our intuitions about them. In our everyday thinking, I claim, we are committed to the view that only some such structures of relations, and only some constituents of those relations, *accurately* and *appropriately* represent their subjects. The phenomenon of aspectual thought and the uses to which we put it, especially in metaphorical communication but more generally as well, reflect this normative commitment.

The task of this chapter, then, is to identify a pattern of thoughts which are associated in such a way that their associations can be significantly restructured without changing what is thought about entirely. In addition, I will argue that these associations are normative and not merely causal: different structures of associations represent the same subject in significantly different ways. In chapter 4, I describe how the application of an aspect brings about the relevant restructuring, and what the effects of this restructuring are.

3.2: *Characterizations*

I will develop Black's suggestion of a 'system of associated commonplaces' in terms of what I call a *characterization* of an individual or kind. A characterization is in the first instance a representation of a particular individual's *character*, or of the character that individuals of a certain kind tend to have. (I'll explain in §3.6.1 under what conditions I think we can have characterizations without characters.) Actions, events, and situations can also be characterized, but when they are, the characterization is derivative upon the characterizations of the (kinds of) individuals participating in them, and the former

characterizations are less rich and complex than those of the individuals on which they are based.[4] Although 'characterization' and 'character' are my own, technical terms, I intend for 'character' to have its roots in the familiar sense of personality, as well as in the perhaps old-fashioned sense of an individual's distinguishing features. The essential defining feature is that having a character requires more than merely possessing an accidentally associated collection of properties. First, those properties must '*hang together*' in various ways. Likewise, a characterization of some thing represents that thing's properties as hanging together in a complex structure. And second, that way of hanging together must make some further properties and not others *appropriate*: as we say, they may be 'in' or 'out of character'.

I think that we do regularly employ characterizations in our dealings with the world around us. We don't just have lists of beliefs about the people and things with which we interact and the situations in which we interact with them. Those beliefs are constellated, along with other attitudes, into a complex structure into which some new beliefs and attitudes but not others can be easily accommodated. And I think the structure of these mental representations, and the constraints on accommodation which they impose, have important effects on the processes of thought in which we normally and naturally engage. I'll employ examples throughout this chapter to suggest how common our use of characterizations is, and I'll discuss their effects on our patterns of thought in chapter 4.

This way of thinking about individuals and kinds, and in turn about their associated actions and situations will, however, systematically come apart from a purely mechanistic view of the world. On such a mechanistic view, there are no real individuals with robust identity-conditions. Rather, there are just more or less orderly piles of material, or more or less dense fields of molecules or energy, which are shaped and re-shaped entirely by certain physical forces. How 'individuals' end up being, on a mechanistic view, is fully determined by the forces and interactions that happen to cross their paths. By contrast, when we think of the world as containing real individuals with characters, we think of those

---

[4] As should become clearer in the next two paragraphs, this is because only individuals have the robust identity-conditions needed for characters.

individuals as reacting to those forces in a particular way: as resisting those changes that threaten their characters, and as embracing those changes that comport with and enhance that character.[5]

We can make the easiest and fullest sense of persons as having characters, but we also have characterizations of many other sorts or individuals as well: animals, artifacts, even in many cases natural inanimate objects. In some of these cases, we may characterize an individual or kind without being committed to the claim that that individual or kind really *has* a character. I don't want the notion of 'characterization' to be tied to character so closely that it rules out these cases, so my account of characterizations should be robust enough to allow them to stand on their own. I'll discuss these possibly 'characterless' characterizations, and the status of judgments about character more generally, in §3.6, after my story about characterizations is on the table.

As we'll see, characterizations are importantly associated with concepts. Most fundamentally, a being who lacked the sophisticated cognitive resources necessary for possessing concepts would also thereby lack the resources for having characterizations. More specifically, particular characterizations are intimately associated with particular concepts. However, this relation must be merely one of association, not of inclusion. Characterizations need to be distinguished from concepts in three key respects: their *structure*, the sort of *commitment* required for inclusion within a characterization, and which *sorts of features* they typically include. I discuss these three factors in sections §3.3, §3.4, and §3.5, respectively. Taken together, these differences make characterizations much more idiosyncratic and contextually malleable than concepts.

Further, while concepts form the basic building blocks of thought, without which no thinking is possible, characterizations are not essential in the same way. Someone who did not characterize the individuals and kinds around him in the way I describe could still think and make claims and other speech acts about those individuals and kinds. I think such a person would have a particularly hard time making sense of the pragmatic nuances of communication and of our aesthetic assessments of the people and things around us, but this would not render him inherently incompetent as a thinker or a speaker. I

---

[5] Cf. Joel Kupperman, *Character*, p. 7.

contrast characterizations, concepts, and words in §3.7. In §3.8, I contrast characterizations with other sorts of mental machinery, such as prototypes, that have been proposed by philosophers and especially psychologists.

3.3: *The Structure of Characterizations*

Characterizations represent individuals as having a set of properties that hang together in a complex structure. This structure in turn varies along at least three interacting but mutually irreducible dimensions: prominence, centrality, and depth. We might think of these as three rather different ways in which a feature might play an important role in our characterization of an individual or kind.

First, as we saw with aspectual perception, some constituent elements are more *prominent* or 'highlighted' than others. When we see the 'old crone/young lass' figure as an old crone, the elements corresponding to the crone's nose and chin are more prominent than the one corresponding to her eye. So too for characterizations: someone's past military service, or pointy incisors, may be quite prominent features in my characterization of him; while those same features may be quite unobtrusive in my characterization of someone else. Similarly, the fact of having been made in China might be a highly prominent feature in my characterization of, say, a Tang funereal figurine but unobtrusive in my characterization of my sneakers.

Prominence is roughly equivalent to salience, at least on some common ways of understanding salience. I will appropriate for my account of prominence the two principal factors that Amos Tversky offers as determining the salience of a given property Φ. First, Tversky says, Φ may be more or less *intense*, in the sense of having a high 'signal-to-noise ratio' or of being obvious. Intensity is exemplified in "the brightness of a light, the loudness of a tone, the saturation of a color, the size of a letter, the frequency of an item, and the clarity of a picture or the vividness of an image."[6] So, for example, the mousy librarian's bright red, racy glasses may rank high in intensity, because they 'stick out' relative to the rest of her more placid demeanor.

---

[6] Amos Tversky, "Features of Similarity," p. 342.

Second, Φ may be a more or less powerful *diagnostic* tool: it may be more or less useful for classifying objects within a given context. Diagnosticity is a more relevant component of salience to the extent that the Φ-ish object is located in a context in which there is a set of clear alternatives to being Φ, and in which being Φ matters. For example, *having a triangular head* would be a highly salient property on a hike where I needed to sort snakes into poisonous and harmless by attending to their heads' shapes. Apart from such a specific context, we do still retain a more general sense for which features are useful for assessing what sort of thing a is. For instance, we tend to treat education, occupation, and where one grew up as relatively important diagnostic features for people.

The second structural dimension of characterizations is *centrality*: what we take to be the feature's importance to determining the *rest* of that individual's character. We saw that in perception, different aspects don't just highlight different features, but also fix certain features and not others as 'going together'. So here, characterizations represent some of an individual's properties as connected to, and in particular as the product or manifestation of, others. For instance, we might interpret someone's adult interactions with strangers, superiors, and small children as manifestations of his having been the childhood victim of a schoolyard bully; or we might take the bullying to be a relatively isolated historical fact. One way of testing for a feature's degree of centrality is how much *else* about the individual's character we think would change if that feature were removed.

The simplest and most obvious analysis of centrality is causal: central features cause many other features. I think something like this may be right. However, that analysis needs to be qualified in at least two important respects. First, the sort of causal connection in play is quite loose and motley: it is often the sort of high-level causation we invoke in explaining action and responsibility, rather than the sort of low-level causation we invoke in explaining physical processes. Indeed, I suspect that we often first think that one feature 'motivates' or 'explains' or 'makes sense of' another, and that this thought in turn generates the faith that *some* causal law or other must explain that connection in *some* way.

The second important qualification is that the causal agent need not be the more central feature. So long as there is some sort of causal connection between two elements, the order of centrality can run either way. We can focus on the causal agent and trace out its manifestations, or we can focus on the result and trace out its various causes. Sometimes two features may be connected by virtue of both being causal manifestations of some other factor, which may or may not itself be known. Which of those elements is more central depends on which one is in turn connected to more other elements. So, for instance, if we contrast two features of Charles's — his being aggressive and his having been bullied as a child — either feature may turn out to be more central, depending on the rest of his character. On the one hand, his aggressiveness may be the more central feature in his character, and in turn in my characterization of him: it may be caused by his having been bullied, and by his high level of testosterone (which also causes his baldness, his moodiness, and his love of fast driving); and it may in turn cause him to be a successful salesman, to get good tables in restaurants, and to lose jobs on a regular basis. Or, on the other hand, Charles' having been bullied may be the more central feature: it have been caused by his slight stutter and his being the poor child of a single mother; and it may in turn cause him to have difficulty looking people in the eye, to challenge authority figures, but also to be solicitous of small children's feelings.

In Figure 3, I have provided a relatively crude, partial representation of a possible characterization of pigs. In a more complex model, we would expect both other local centers of connectedness, and also more subtle distinctions in prominence. Prominence is represented by the heaviness of the framing boxes, centrality by the arrows. It is important to note here that prominence and centrality are often correlated, but that neither can be reduced to the other. Differences in centrality produce some differences in prominence, but not all prominent features are central, or vice versa. For instance, in our characterization of pigs, our image of their fat snouts is a prominent but not a central feature. Similarly, it might be a quite prominent feature of Bob's character that he once saved two children from a burning building, but this might not be connected to many further features of his personality. Conversely, Charles's childhood experience of being bullied might not itself be all that

prominent, but might still play a central role in determining his other character traits. In Figures 4a and

4b, I've represented two rough, schematic characterizations, illustrating the ways in which taking

FIGURE 3 HERE

FIGURES 4A and 4B HERE

different features to be central can rewrite the structural relations of centrality and prominence among other features in that characterization.

So long as we remain within characterizations themselves — *our* patterns of thought *about* individuals and kinds — prominence and centrality are, I think, on the same footing. However, there is an important difference between them at the level of characters — of how individuals really are. Prominence is highly dependent on observers' interests in the particular context of application; centrality is not observer-relative in the same way. While we can talk about a feature's being prominent relative to many different observers' interests and across many different contexts, it doesn't make sense to claim that a certain feature *just is* prominent in an individual's character, *tout court*. By contrast, it does make sense to maintain that a feature is central to determining an individual's character regardless of our views on and interests in the matter. One might, for instance, discover that one was wrong in what one took to be central to someone or something's character. If we contrast the two characterizations of John in Figures 4a and 4b, we can imagine two people arguing about what really 'makes him tick', where each represents John in one of the two ways illustrated. This disagreement is a genuine one, about how John really is. Because centrality is a feature of characters themselves, judgments about centrality have an objective status that judgments about prominence lack.

Centrality in this sense thus needs to be clearly distinguished from *theoretical* centrality. This latter notion captures the idea that some beliefs are highly confirmed, and so located at the center of our net of beliefs, in Quine's image.[7] As a result, those beliefs are highly insulated, though not in principle immune, from revision. So, for instance, the shift from believing that the sun revolved around the earth to believing that the earth revolved around the sun represented a shift of great theoretical centrality. But it is unlikely to have been a shift of great centrality in my sense. This is so for at least two reasons: revolving around the sun doesn't explain or cause many other features in our characterization of the earth; and as I said above, our everyday characterization of the earth is primarily focused on our ordinary experiences rather than on current scientific theory.

---

[7] "Two Dogmas of Empiricism."

Finally, in addition to being prominent and central, some features run *deep* in one's character.[8] Deeply rooted elements are highly impervious to change. Such features need not be prominent: they may reveal themselves only occasionally, and even then not very obviously. Nor need they be central: they may be quite isolated from other elements of one's character. Perhaps someone's bigotry, or fanaticism, might be deep in this way. The person may have taken care to excise bigoted ways of talking, and even of thinking, from most of his life; but in certain situations he may still feel a violent revulsion at the presence of someone whose kind he disdains. Depth won't play nearly as significant a role as prominence or centrality in my account of aspectual thought and metaphor, but I think it should be included because it plays a significant role in our understanding of people. (I'm less convinced of its relevance to the characterizations of beings without psychologies.)

Now, given these three distinct dimensions of structure, it should be clear how much it would distort our characterizations to reduce them to checklists of properties believed to be possessed, *even* if those properties were hierarchically ordered. With characters and characterizations, the whole is more than the sum of the parts. Two individuals who possess nearly the same set of properties may still differ markedly in character, because those elements hang together in such different structures. *A fortiori*, two thinkers may have very different characterizations of the same individual without disagreeing about which properties it possesses, because they may structure those properties quite differently. When those differences merely involve prominence, two thinkers will notice or cite different features of the subject under consideration. As a result, they will notice different similarities between that subject and others. When those merely structural differences in characterizations turn on centrality or depth, as they often will, the characterizations will objectively differ in how they represent their subjects as being.

3.4: *The Relation of Characterizations: Fittingness*

The next important question about characters and characterizations is which features go to make up an individual's or kind's character, and more importantly for our topic here, which features should be

---

[8] I owe this point to Richard Wollheim.

included in our characterization of it. ==Obviously, most of the elements in our characterizations must be properties that we think the individuals characterized actually possess, and vice versa.== We'd lose our grip on what an individual's character was if we strayed too far from actually-possessed features. However, commitment to actual possession is neither necessary nor sufficient for including an element within a characterization.

On the one hand, possession is not sufficient for inclusion. We rightly dismiss some things we know to be true of an individual as being out of character. We may do this because we think those features are merely incidental or irrelevant. For instance, not everything I know and know to be known about pigs is part of my characterization of pigs. The fact that they are mammals isn't particularly relevant, partly because it's not a particularly distinctive feature of them. And the facts that they are non-ruminants, and that there are eight major breeds of pig in the United States, are simply irrelevant to our common characterization.

But we may also dismiss a feature as out of character because it conflicts with the rest of our characterization. So, the fact that pigs are actually quite clean animals conflicts with our common characterization of them. The fact that a possessed feature conflicts with our characterization may be a sign that our characterization is inaccurate: maybe we've got pigs all wrong, because their being dirty is so central and prominent in our characterization. But sometimes a feature really *is* out of character. The sweet, kind librarian may one day snap and refuse to help an addled patron find just the right book on begonias. Later, she herself may dismiss her action, saying something like 'I just don't know what came over me; that's not who I am.'

Animals too can do things which are out of character, for them as individuals and for their species. One lioness in Africa has apparently been isolating baby gazelles from their mothers and gently coddling them.[9] (The problem is that she doesn't feed them, and so they eventually starve to death. Then, but only then, she eats them.) Because this particular lioness has done this repeatedly, we may

---

[9] "5 Little Oryxes and the Big Bad Lioness of Kenya," *New York Times* (October 12, 2002).

need to include this behavior within *her* character. But we needn't revise our characterizations of

lionesses in general just because this one indulges in this strange behavior: we can dismiss these actions

of hers, or her character more generally, as a fluke. Even artifacts can possess properties which are out of

character. We often criticize artworks for containing elements which don't fit, given the individual work

of art or the overall style of the artist. We also criticize artifacts other than artworks in the same way: the

fabric in which a chair has been reupholstered, for instance, may not comport with the rest of its

character.

So it should be pretty clear that not all possessed properties need be part of an individual's or

kind's character, nor of our characterization of it. ==The converse claim, that actual possession is not==

==*necessary* for inclusion within a characterization, is more bizarre but also, I think, true.== One reason for

accepting this claim is Black's point about 'associated commonplaces' above. Some 'commonplaces'

may be "downright mistakes" from my point of view. They may nonetheless still be included in my

'system' of thinking about the thing in question, because I know that people generally think that the

individual or kind possesses the property. For instance, I may know that gorillas are generally gentle

creatures, and I may even know that my interlocutor knows this, but it may still be mutual knowledge

between us that many other people think that gorillas are brutish and violent, and so we may each include

that feature within our characterizations of gorillas. Similarly, I may know that pigs are in fact clean

animals, but not be able to reconcile this with the rest of my pig-characterization, and so I may continue

to include being dirty within my characterization of them as untrue but fitting.

Further, we have characterizations of fictional and mythical 'individuals' and kinds; and we

might think that there can be no possibility of actual possession here.[10] Notice that our intuitions about

which features should be included in our characterizations of these individuals and kinds often goes

beyond what is actually stipulated in the fiction or myth. We may have strong intuitions about how

Sherlock Holmes speaks, or walks, or examines a piece of evidence, even though Conan Doyle never

---

[10] Though, depending on one's metaphysical views, unreal characters may actually possess properties.

explicitly addresses these matters. We think that some of these intuitions are correct and others wrong, and we argue with each other about who has a better understanding of Holmes' character.

In explaining metaphorical comprehension in particular, most of the unpossessed properties included within characterizations will be of one of these two sorts: 'downright mistakes' or properties of fictional or mythical individuals or kinds. However, I want to maintain the stronger claim that unpossessed properties can sometimes rightly be included in an individual or kind's character, and so in our characterization of it, *even* when no such mistake or fiction is involved. I think this point is especially important for understanding our engagement with fictional and historical narratives, and for aesthetic appreciation more generally. I will say that our characterization of an individual or kind includes all and only those features we represent as *fitting* its character. We dismiss possessed features as out of character when they aren't fitting; and we include features which aren't possessed within our characterizations when they are fitting.

Let me bring out the point about unpossessed but fitting features by example. Sometimes we feel that an individual, especially but not exclusively a person, so fully exemplifies a certain way of being that it would be fitting for that individual to have a certain other property, regardless of whether or not it actually has it. For instance, Bill may be such a quarterback sort of guy that it would be fitting for him to be, or to have been, a quarterback, even though in fact he never was one. Knowing that he never was a quarterback may do nothing to sway my assessment of its fittingness for him: I may dismiss this fact by saying something like "Well, he *should* have been a quarterback." An apocryphal tale about a famous personage may persist despite widely available proof of its falsity, because it is so fitting. The people who repeat the story may say be relatively unconcerned about its actual truth, and more concerned about what it reveals about the personage's overall character.

Features even can be fitting for someone or something when they aren't viable possibilities for them. It might be fitting for a young man with a certain pent-up fervor to throw himself into religious martyrdom, but he may live in a culture and an age in which this is impossible. Or it might be fitting for an aging woman to reign as the neighborhood matriarch over afternoon teas, even though she has been

forced by immigration and poverty into working as a janitor. In a case like this, I will refuse to include the woman's poverty-stricken state and occupation within her character, because I think they don't fit it.

Presumably, facts about the fittingness of unpossessed properties must eventually supervene on more complex collections of facts: about properties that the individual or kind characterized does actually possess, and about the more general social and natural environment in which that individual or kind is located. After all, intuitions about fittingness respond to something about how the individual in question actually *is*. However, it's not very informative simply to claim that there is some such supervenience relation; and I suspect that not much more can be said. Facts about the fittingness of properties are not reducible to facts about actual possession in any simple or systematic way.

It might seem that facts about fittingness should at least be analyzable in dispositional terms: although *a* hasn't in fact been placed in circumstances appropriate for it to be *F*, if it were to be placed in such circumstances, then it would be *F*. This sort of analysis seems like it might take care of the examples above, and of others as well: that Alice is such that it would be fitting for her to be Queen of an exotic Polynesian island; or that Carol is such that it would be fitting for her to join a religious cult. There's something about how each of them now is which can only be captured by saying that this is what they *would* do if the circumstances were right.

However, a dispositional analysis can't work in general, because there is often no clear sense in which the relevant circumstances are themselves possible. For instance, we might have the intuition that it would be fitting for a grand and gnarled old oak to speak in a low, resonant voice, but this is relatively high on the scale of impossibilities; on many selection-restriction semantic theories, it's simply nonsensical. Similarly, and less outlandishly, we often think that a historical figure had a character that makes it fitting for him or her to be portrayed by this actor but not that one, for that actor to carry himself in a certain way, for a certain sort of theme music to be played upon his entering the scene, and so on. But in thinking all this, we are clearly not ascribing to the historical figure himself the dispositions to behave in certain ways had *he* been placed in a certain situation. Nor does any other such disposition seem forthcoming.

It may also be fitting for an individual *a* to have had some (non-actual) event in its past, which would have contributed toward forming *a*'s current (actual) character. But if I claim that it's *now* fitting for there to have been such a past event, I'm not claiming that *a* actually *had* the relevant disposition at that past time but just happened not to encounter the right circumstances for manifesting it. So, for instance, Diana may now be such that it would be fitting if a fiancé had abandoned her at the altar many years ago — she has just *that* sort of dark suspicion and worldly disappointment. No such event actually occurred, though, and I know this. What actual disposition could I now be attributing to the Diana of that time, many years ago? Certainly not the disposition to be darkly suspicious: she was bright-eyed and innocent then, and I may know this and accept it as fitting. I am trying to capture the way someone who was like *that* then could come to be like *this* now.

Finally, it looks like intuitions and facts about fittingness can't be reduced to intuitions and facts about possibility more generally. So, for instance, Fred's being a father, and being rather colorful and theatrical, are much more fitting for and central to both his character and my characterization of him than his genetic code is, or which parents he had. But the latter properties play a much greater role in determining the relative nearness of possible worlds with respect to him. So too, I may think that *if* Bob were to be a non-human animal, it would be most fitting for him to be a bull. But this possibility is in fact much more remote than that of his being a chimpanzee, given the genetic makeups of these three species. Similarly, it would require just as much metaphysical voodoo to turn the Campanile into a mammoth jutting icicle as it would to turn it into a placid lake, since a lake and an icicle are made of the same basic stuff. But being an icicle is much more fitting for the Campanile than being a lake is.

So, although intuitions about fittingness overlap and intersect with intuitions involving more familiar metaphysical notions, I'm skeptical that the former can be brought within the fold of the latter. We do nonetheless still have intuitions about fittingness. In the coming chapters, in my account of the sort of aspectual thought relevant to metaphorical comprehension, fittingness will primarily serve as a shorthand for that subset of actually possessed properties which comport with the individual or kind's overall character, plus any unpossessed properties which are commonly thought to be part of its character.

98

But because I think that the same basic patterns of thought underlie our engagement with fictional and historical narratives, and because I think these narratives do sometimes depend on unpossessed but fitting properties, I make the further (and stranger) implications of fittingness explicit here.

Before leaving the discussion of fittingness, I want to mention a distinction between two ways in which one might entertain the thought that some feature, *F*, is fitting for an individual *a,* or kind *A*. On the one hand, one can entertain what I'll call a *general* thought that *F* is fitting for *a*. The judgment is 'general' in the sense that it affirms merely that there is *some* way of fitting *F* into *a*'s character, without staking any sort of commitment as to just which way that is. On the other hand, one might entertain the *singular* thought that *F* fits *a*'s character *in a particular way*. Doing this requires actually fitting *F* into one's own characterization of *a*. This in turn means locating *F* with some rough measure of prominence, centrality and depth within the structure of one's characterization of *a*, and getting some sense for how *F* connects up with the other elements in that structure. Adding a new feature to one's existing characterization in this way will usually also force alterations in the rest of that characterization, by motivating some new structural relations between already established features and undermining some others. This distinction will become important for us in chapter 4, because applying an aspect typically requires entertaining the singular thought as a first step.

3.5: *The Elements of Characterizations*

In this chapter, I am developing my own, more explicit and accurate account of the 'system of associated commonplaces' which Black claims are essential to comprehending metaphor, and which I think serve as the material on which aspectual thought operates more generally. I call the relevant 'system' a *characterization* of the relevant individual or kind. In the last two sections, I have argued that characterizations have a complex, three-dimensional *structure*, and that they include all and only those features we represent as *fitting* the character of the individual or kind characterized. In this section, I discuss the sorts of features which can be included.

Black claims that only 'commonplaces' are included in his 'systems'. However, this doesn't provide a workable criterion for determining which sorts of features can be included within characterizations, and not merely because the notion of a commonplace is itself so vague. First, we don't just have characterizations of kinds like *man* and *pig*, but also of individuals like *George W. Bush* and *my colleague Charles*. It's much less clear what would count as a 'commonplace' about an individual. Further, and more seriously, we don't just include in characterizations features which the individual or kind is generally known to possess, but also quite specific knowledge about them.

When we get to metaphorical communication specifically, something like Black's 'commonplaces' will come to be more relevant: there must be something like mutual knowledge between the conversational participants in order for communication to succeed. Even there, though, the relevant sort of commonality will be much more nuanced and particular than talk of 'commonplaces' suggests. In a specific conversational context, all sorts of highly specific, special assumptions may be common to the conversational participants and relevant to metaphorical interpretation. Black himself points this out: "Metaphors can be supported," he says, "by specially constructed systems of implication as well as by accepted commonplaces; they can be made to measure and need not be reach-me-downs."[11] He still seems to assume, though, that in such a case the system of implications will be explicitly constructed by an author in advance of the metaphor itself, and this need not be the case. More importantly, characterizations and aspectual thought don't just play important roles in communication, but also in individual contemplation; and here nothing like the commonality requirement obtains. An account of characterizations and of aspectual thought which could only explain their roles in metaphorical communication would provide a distorted picture of the phenomena themselves.

As we've already seen, many of the constituents of character are ordinary properties like *being intelligent*, *being dirty*, or *being aggressive*. Correspondingly, many of the constituents of characterizations are straightforward general concepts of those properties. In addition, though, characterizations can and typically do include two sorts of elements that have until recently been

---

[11] "Metaphor," p. 43.

relatively neglected in philosophical theories of mind: highly experiential features for which we lack standing terms in a public language, and the fittingness of certain affective attitudes for the subject characterized. Philosophical theories of mind have tended to neglect these sorts of features because it can be difficult, and sometimes impossible, to find general literal expressions denoting them. Because we are often driven to speak metaphorically when we want to refer to such features, it is especially important that I bring attention to them here.

3.5.1: *Experiential features*

Characters can include highly specific properties for which we lack standing terms in the language, such as being *that* shade of scarlet, or (returning to Figure 3) as making *this* sort of snorting sound. Many philosophers now seem willing to acknowledge that we can form demonstrative, experientially-based concepts of such properties,[12] but this is a comparatively recent shift, and one that is not yet fully settled. Christopher Peacocke, for instance, has sometimes ~~apparently~~ restricted~~s~~ concepts to constituents of thought that correspond to standing terms in a public language.[13] I think that this restriction is wrong: we *can* form demonstrative concepts for properties to which we cannot refer using public linguistic terms in their literal meanings.

One reason for restricting the range of possible concepts to those that are expressible in a public language might be to respect what is sometimes called the *general* or *abstractive* nature of concepts. Because concepts essentially classify particular experiences together, essentially particular reference to features cannot be genuinely conceptual.[14] In order to respect this 'abstractive' quality — to ensure the "necessary distance from what would determine them to be true" — John McDowell, for one, requires that a demonstrative concept must be available in principle for some time "beyond the duration of the

---

[12] *Cf.* e.g. David Kaplan, 'Quantifying In,' p. 142; Wittgenstein, *Philosophical Investigations* §78; McDowell, *Mind and World*, chapter 3. ~~,~~

[13] See "Scenarios, Concepts, and Perception." Philip Pettit (in conversation) has articulated the same view.

[14] See Sean Kelly, "The Non-Conceptual Content of Perceptual Experience: Situation Dependence and Fineness of Grain." ~~Thus,~~ Alexander Nehamas ~~holds~~ has said (in conversation) that the propositional is limited to what is abstractly transmissible, while the non-propositional corresponds to what one must 'work out for oneself'. ~~-~~

experience itself."[15] I'm not sure just how McDowell wants this requirement to be applied, but I suspect that it sets too high a standard for concept-possession.

That is, if the following judgmental and inferential conditions are met, then I think we would willingly and correctly ascribe the relevant concept to the relevant thinker. The thinker must have a robust ability to identify the property when confronted with an instance of it; she must be able to make sense of some things and not others possessing the property; she must be able to draw certain conclusions from the thought that something does possess that property; and she must understand that she might be wrong in her judgment about whether a particular thing does possess it.[16] But all of this might be true even while the thinker must rely on a sample of the relevant property, or on some other procedure for directly comparing the new putative instantiation of the quality to the original, definitive one. In such a case, the thinker's conceptual capacity is in some sense essentially 'anchored' to the original instance, but the capacity plays a sufficiently rich role in her thinking that it should still be counted as a genuine concept.

A related source of resistance to demonstrative concepts might stem from Wittgensteinian worries about the possibility of a private language. One might think that because demonstrative concepts rely essentially on personal experience without being tied to standing public terms, they can't allow for public agreement about the criteria for their correct application, and so can't be genuinely normative in the way that seems to be constitutive of concepts. We can address this worry by requiring that some sortal concept be included within the demonstrative concept, so that there are no genuinely 'bare' demonstrative concepts, but only elliptical ones. I am inclined to accept this requirement. However, the sorts of thoughts about features which are relevant for characterizations will generally meet this condition: one thinks of the relevant property as, say, *this tone of voice* or *that shape*.

In some such cases, the accompanying general concept may still not fully fix just which qualities of the tone, or shape, or other property must be present in order for *another* voice or shape or whatever to

---

[15] *Mind and World*, p. 57.
[16] Cf. Gareth Evans, *Varieties of Reference*, Wilfred Sellars, *Empiricism and the Philosophy of Mind*.

count as relevantly the same. This is especially true when we come to thoughts about quasi-aesthetic features like gestures and styles. The paradigm case of a demonstrative concept seems to be a color concept: what's expressed by suitable utterances of '*that* shade of scarlet', or more simply, '*that* color'. Even here, there is still some vagueness about just how to draw the concept's boundaries, but we can at least be sure of the dimensions along which that vagueness ranges: at most, colors vary in hue, saturation, and brightness.[17] By contrast, general concepts like *gesture* and *style* aren't associated with any clear dimensions of variation. As a result, the criteria of correctness for compound demonstrative concepts involving them will be quite heavily context-dependent.

This point is important because these sorts of features play a crucial role in our characterizations of many individuals and kinds. We typically associate certain styles, certain ways of moving or carrying oneself, or of being responded to by others, with particular individuals and kinds of individuals. (Think of how one would imitate or recognize the Queen of England, or a caged tiger, for instance; and how it is fitting for one to act in relation to them.) And we think that those gestures and styles reveal something important and distinctive about those individuals — indeed, these features are often *more* revealing than their more explicitly goal-directed behavior. But we often also find it very difficult to describe those features using just literal language, without ostending the feature itself.

I still think demonstrative concepts involving general concepts like *gesture* or *style* set sufficiently tight constraints on re-instantiation that we should be entitled to "talk about 'right'" here.[18] At the very least, these demonstrative concepts fix as much of a "criterion of correctness" as corresponding *non*-demonstrative concepts and terms do. In the abstract, without contextual specification, we might be quite unsure whether someone exemplifies 'the style of Coco Chanel' or 'the fluid sweeping movements of Martha Graham'.[19] But we usually think that these phrases are meaningful,

---

[17] Though Merleau-Ponty, for one, argues for a much more robust context-dependence even for colors: the red of a steel ball is not the same red as the "wooly red" of a carpet, he claims, even if the hue, brightness, and saturation are the same (*Phenomenology of Perception*, p. 5); see also Kelly, "The Non-Conceptual Content of Perceptual Experience."

[18] *Philosophical Investigations*, §258.

[19] Within the context of choreography, or of practicing a dramatic routine, there might be determinate, though not

and fix genuine truth-conditions. Further, this indeterminacy is shared to a significant degree by thought and language more generally. We may be similarly unsure, for instance, whether a given object instantiates the concepts *green* or *table* even though we know all the relevant facts about the object, about what it means to be green, and about tables.[20] Thus, the indeterminacy we find here shouldn't call into question the legitimacy of demonstrative concepts specifically, or even of demonstrative concepts that refer to quasi-aesthetic features like gestures and styles.

3.5.2: *Affective Attitudes*

In addition to experiential concepts of highly specific features, characterizations often include affective attitudes, emotions, and moods; these attitudes often play important roles in those characterizations as central, prominent, or deep elements. For these attitudes, I think it is much less clear that we can always represent these elements in a fully conceptual or propositional way. To the extent that this is so, and that these features also play an important role in metaphorical comprehension specifically, we have another source for the intuition that such comprehension involves a sort of non-propositional understanding.

Affective attitudes can play two sorts of roles within a characterization. On the one hand, we may think it is fitting for the individual or kind characterized itself to have a certain attitude, emotion or mood. So, we might impute a feeling of wounded pride to a caged lion or to a downsized CEO, or a feeling of defiant glee to the arsonist standing before his latest crime. On the other hand, we may think it is fitting for *us* to have certain feelings in our interactions with that individual or kind. So, I might think it would be fitting to feel rage toward that arsonist, to feel awe before Mother Teresa, to feel piety when standing in a grove of aspens, to feel simultaneous curiosity and revulsion before an artwork by Damien

---

explicitly stateable, conditions of satisfaction for a particular gesture or posture. But even here there are difficulties: apparently the history of dance is fraught with uncertainty over whether *the same moves* are being performed now as were performed previously.

[20] John Searle, Charles Travis, and others have argued on these Wittgensteinian grounds that all language requires additional appeal to background forms of life in order to fix genuine conditions of satisfaction. See e.g. Searle, *Intentionality*, "Consciousness, Intentionality, and the Background"; Travis, "On Constraints of Generality," *Unshadowed Thought*.

Hirst, or to feel raucous excitement in sports stadiums. One need not actually feel the emotion in question for it to be included in one's characterization, so long as one takes that emotion to be fitting.[21] However, when one is actually confronted with the relevant individual or kind, and one thinks that a certain emotion is fitting for it, then one will naturally tend to feel that emotion, at least to some degree. If one doesn't feel that way in its presence but still maintains that this is the fitting way to feel, then one needs some special explanation of the disparity.

Now, can these emotions and moods be represented in a fully propositional way? We obviously do have a battery of concepts for affective attitudes: the concepts *fear*, *joy*, *loathing*, *pride*, and so on. Many philosophers would agree that such sentences as

(3) John is happy to see Jane,

which are partly about affective attitudes and emotions, express propositions which can then serve as the contents for propositional attitudes and be evaluated for truth and falsity in the usual ways. And it then seems in particular that those propositions can be embedded under an 'it is fitting that' operator to form more complex propositions. However, there are two putative reasons for thinking that this can't provide a satisfactory, fully propositional representation of the role that affective attitudes play in characterizations. I find the second reason, but not the first, to be compelling.

The first putative reason is that in order meaningfully to ascribe an attitude, emotion, or mood to someone, we must at least have some idea *which* attitude we are imputing.[22] Sometimes, the emotions and attitudes that we want to impute to others are highly specific; when this is so, demonstratives may

---

[21] Cf. Richard Wollheim on expressive properties in *Art and Its Objects*, pp. 56-61, and in "Correspondence, Projective Properties, and Expression."

[22] Cf. Evans, *Varieties of Reference*, ch. 6. One might be able to *utter* a sentence with a certain content without meeting this requirement, but then one wouldn't fully understand what one was saying, or claiming. The 'knowing which' requirement can be fleshed out in many ways, which may differ significantly. Some philosophers worry that any such requirement is too stringent. Cf. here Kaplan's example (in "Dthat") of mistakenly ostending a picture of Spiro Agnew rather than one of Carnap, as he intends; the intuition is supposed to be that Kaplan thereby makes a claim about Agnew despite his intention. Kaplan himself later says this is an unusual case ("Afterthoughts," p. 582, fn. 34). My own intuition is that his referential intention should be analyzed as 'the man depicted in *that* picture,' and that Kaplan does know *which* picture he is ostending. For reasons which will become apparent in the next paragraphs, I think this is the only sort of attribution one can make of a highly specific emotion (e.g., 'the sort of fear that *he's* feeling') unless one has some direct experience with what it's like to have that emotion. I also think that the 'tracking' requirement which is often taken to underwrite the 'knowing which' requirement is more easily satisfied for demonstrative reference to objects than for demonstrative reference to qualities.

again provide the only means for referring to them. But, one might argue, in order to know *which* highly specific attitude we are thinking about, we must ourselves have experienced that emotion or attitude. Thus, for instance, we can impute the *general* emotions of happiness or romantic love to someone without ever having experienced the relevant emotions ourselves. We can do this because happiness and romantic love play sufficiently rich roles in our common public lives: we know what sorts of things people do when they experience these emotions. But we may not in the same way be able to understand, and so may not be able to impute to others, the particular emotion that someone feels in a highly specific situation: say, the particular combination of admiration, sympathy, and loathing which captives sometimes feel toward their kidnappers.

This requirement of being acquainted with the attitude ascribed might then seem to introduce a non-propositional element into the ascription itself. David Novitz argues for something very like this claim in his discussion of fiction:

> Emphatic beliefs and knowledge of this sort [about "what it feels like to be in certain complex and demanding situations"] are derived from, and have to be explained in terms of, our awareness or experience of what is sometimes termed a direct object. A feature of such awareness is that it is irreducibly non-propositional, in the sense that it cannot be captured or adequately conveyed in linguistic descriptions. For instance, no matter how precise and vivid my descriptions are, they will never acquaint you with my feelings as an orphan, or with the strains of Mozart's *Te Deum*, K. 141, or with the anguish of a moral dilemma... Our imaginative involvement in fiction allows us to respond emotionally or feelingly to the tribulations and triumphs of creatures of fiction, and so to acquire such non-propositional beliefs.[23]

Now, although I agree with Novitz that certain beliefs about how others feel require 'experience of a direct object', I don't think it follows that those beliefs themselves are therefore "irreducibly non-propositional." The situation is precisely parallel to that of demonstrative reference to highly specific properties like color: in both cases, the requirement that one have experienced the relevant property — the color or emotion — is merely a *background condition* on one's being able to entertain the relevant proposition. The proposition itself may indeed include what is ostended as a constituent; but this just implies that it is a singular proposition, not that it is not a proposition at all. Thus, Novitz's claim about the non-propositionality of these emotionally complex beliefs depends upon ignoring or significantly

---

[23] *Knowledge, Fiction, and Imagination*, p. 120.

distorting the role that demonstrative concepts play in thought, and in turn the role that demonstrative terms play in communication.

The second fundamental difficulty with Novitz's claim is that, as Novitz is precisely concerned to point out, we *can* acquire some sort of acquaintance with the relevant affective attitudes 'indirectly', through imaginative engagement. In order to engage fully with fictional or historical narratives, we do need an initial experiential base from which we can extrapolate to an understanding of the situation described. But given such a base, we can use our powers of imagination and empathy to understand situations and characters that we've never directly encountered. This is precisely why reading fiction can be such a powerful — and potentially morally educating — experience.[24] So, although some sort of 'acquaintance' is required in order to gain knowledge of new affective attitudes, we can't interpret Novitz's invocation of 'acquaintance with a direct object' in such a strong fashion that it rules out imaginative projection of these sorts. What Novitz's discussion does bring out, though, is the quite central role that imagination plays in our characterizations — especially in relation to fiction and historical narratives, but also more generally.

The second argument for the claim that affective attitudes introduce something essentially non-propositional into characterizations is, I think, more compelling. This argument turns on the fact that we don't just impute attitudes and emotions to those we characterize; we also find it fitting or appropriate to *respond* emotionally to them in certain ways. Certain emotional and affective responses, and not others, are normatively apt in relation to certain situations and individuals. When we do think that it's fitting for us to respond in these ways, what's inside the scope of the 'fitting' operator is the emotion or mood itself. And these emotions and moods are neither themselves propositions nor even necessarily attitudes toward propositions.

Now, it is true that we can formulate these claims or intuitions as follows: it is fitting *that* one feel *x* in relation to *a*. However, I think this mischaracterizes the initial attribution of fittingness in two ways. First, it carries at least the suggestion of veridicality, that one really *does* feel *x*, and as I said above, this

---

[24] Cf. Martha Nussbaum, *Love's Knowledge.*

need not be so. Second, it covers over the role that the emotion itself naturally plays in the characterization. It's not just the *fact* of there being such a feeling, but the feeling *itself* which is fitting. And in turn, what is involved in having that particular emotion or mood often cannot be adequately analyzed by specifying an attitude toward a determinate proposition which would 'satisfy' the attitude if it accurately represented the actual world.

Nonetheless, affective attitudes, emotions, and even moods are all undeniably world-oriented and world-representing. As Allan Gibbard says, "An emotion...involves a special way of experiencing one's world, a way that will be difficult to express and perhaps can only be whistled."[25] Gibbard's is an especially apt description of moods: the world as a whole looks different when I am ebulliently joyful than when I am anxious and afraid. We might add that emotions and affective attitudes typically involve a special way of experiencing and representing particular individuals *within* one's world.

Emotions and moods are important in our characterizations, not just because they are common constituents of them, but also because they imbue them with vital psychological force. In turn, they are an important source of aspects' and metaphors' much vaunted and maligned powers to cast something in a powerful light. Affective attitudes are so influential in our thinking in part because they involve the rich experiential component pointed to by Novitz and by Walton. But emotions and especially moods are also influential because they are *permeant*: as Wollheim says, they "permeate or infiltrate adjacent or resembling mental states."[26] So, for instance, my joy at my friend's marriage might lead me to find the landscape in which the wedding is held particularly lush and inviting. In particular, affective attitudes typically permeate the *other* represented features in a characterization. They thereby alter the characterization's overall structure, by highlighting just those features which fit with it. Thus, my admiration for Bob's bravery in saving two children from a burning building is likely to render his quiet calm and general competence more prominent, because these qualities fit so well with my attitude.

---

[25] *Wise Choices, Apt Feelings: A Theory of Normative Judgment,* p. 131. Often, we may be able to cite the *causes* of our emotions and moods, but not to describe the particular representational import of the emotions and moods themselves.
[26] Wollheim, *On the Emotions,* p. 75.

Because they are both permeant and experiential, affective attitudes make our characterizations as a whole more vivid.

3.6: *Judgments about Fitting and Character*

At this point I have laid out my account of characterizations, the 'system of associated commonplaces' which I claim underlies aspectual thought in general and metaphorical comprehension in particular. To summarize:

- characterizations are complex, structured representations;
- they represent just those features we think it is *fitting* for an individual or kind of individual, whether or not we believe that those features are actually possessed;
- those features can include both general and highly specific, experientially represented properties and affective attitudes;
- those features are structured according to their *prominence*, *centrality*, and *depth*.
- The character of an individual just is the structured collection of properties and attitudes that are fitting for it, where some elements in that collection are more central or deep than others.[27]

In this section, I explicitly address two issues about the objectivity of characterizations and the normativity of judgments about character. First, I ask whether we really are committed to the metaphysical claim that lots of things have characters. Second, I argue that two thinkers with differing characterizations can genuinely disagree about how the individual or kind characterized *is*, and I discuss how such disagreement can be resolved. In the next two sections, I contrast characterization with other cognitive machinery: with concepts and words, and with more psychological notions like conceptions and prototypes.

3.6.1: *On The Oddness of Attributing Characters*

Characterizations are complex patterns of thought about individuals and kinds of individuals, and derivatively about the actions they undertake, the relations in which they are embedded, and the events they undergo. They are also supposed to be fairly common patterns of thought. Indeed, if they are to provide the basic materials for explaining metaphorical comprehension, and given that the vast majority

---

[27] Prominence, I said, was a function of the context and interests of the characterizers, and so doesn't apply to the characters themselves *per se*, but only to our characterizations of them.

of our words can be used metaphorically, then we will need to be able to associate characterizations with *most* of our concepts and words. But I've also said that characterizing something is attributing a character to it; and having a character sounds an awful lot like having a personality. This seems to commit me to the claim that we are all operating with a highly animistic, personified worldview. And surely, the thought goes, we left *that* behind several centuries ago. Do we really believe that all the things we characterize have characters?

First, thinking of some thing as having a character does not require personifying it. It only requires thinking of it as more than a mere accidental collection of properties. When something has a character, its properties hang together in a complex structure into which some further properties fit and others don't. We do, I think, ordinarily treat many things as meeting this requirement. We do think that some properties are more or less essential to making a thing what it is. We think that those more essential properties determine a lot of what else is true about the thing, and that some changes in which properties it possesses, and not others, would accord with the thing's overall identity. Further, although our intuitions about character do have their antecedents in intuitions about personalities, not all characters need be even vaguely humanoid: we also have a whole range of intuitions about mechanical and animalic characters, which we attribute to humans and to other objects.

Second, I don't think we folk have left this worldview as far behind as we philosophers sometimes like to think. There are genuine, deep conflicts between a fully scientific, mechanistic perspective on the one hand, and on the other not just folk psychology but folk understanding and explanation more generally. The story I've been telling about characterizations is a piece of Strawsonian descriptive metaphysics, an account of the metaphysics embodied in our ordinary ways of speaking; and I maintain that we do often talk and think as I've described. It is oddest to talk about inanimate objects as having characters. Even here, though, I think the notion of character finds genuine application in our ordinary thought. We typically *are* committed to the idea that artifacts — especially artworks — have characters. In these cases, the characters may supervene in large part on the character and intentions of

the sort of person who would create such a thing and on the responses they tend to arouse in people interacting with them. But they are characters nonetheless.

Third, characterizations and characters can be more or less rich and complexly structured. In order to motivate the claims that we do have characterizations, and that we need to appeal to the distinctions I've been making in order to describe them accurately, I have cited relatively complex characterizations as examples. People in particular have much more rich and complex characters than other individuals, and we have the most nuanced interest in characterizing them. Not all characterizations and characters exhibit this much complexity. Indeed, in §3.8.2 I will suggest that we can profitably treat stereotypes as relatively coarse characterizations. What is important is that we can treat these simpler, less rich and complex cases as versions of the same basic phenomenon, in which the same structural dimensions and sorts of features are still represented.

Finally, as I mentioned in §3.2, I do want to make room for some characterizations of (kinds of) things which the thinker does *not* treat as actually having characters. Certain sorts of scientific understanding are the most plausible candidates here. For instance, theoretical physicists develop nuanced and robust ways of thinking about sub-atomic particles and their interactions, but they are only interested in the truth about which properties the particles actually possess. If their working understanding meets the two essential conditions I outlined above, though, then these ways of thinking will still count as characterizations — even if the scientists are not committed to there really being a corresponding character. First, their pattern of thought must manifest a complex structure of prominence and centrality, rather than simply remaining a set or list of properties ascribed. The connections which determine centrality will here be more straightforwardly causal, in a basic and clear-cut sense of causality, than in many other characterizations.

Second, their pattern of thought must incorporate intuitions about which properties fit and which don't. This is a more dubious assumption, but it still seems likely. These scientists might have and rely on such intuitions in their working investigations without actually being committed to them as part of their ultimate theory. For instance, certain sorts of hypotheses might just not 'feel right' to them,

although they are acknowledged to represent coherent, even equally metaphysically likely, possibilities. The fact that physicists rely so heavily on just those hypotheses that can be formulated in elegant mathematical terms might be an example of this sort of selection in terms of fittingness. Most physicists acknowledge that inelegant hypotheses could possibly be true, but they feel it would be more satisfying if the elegant ones worked out.

Still, one may well remain skeptical, especially of the notion of fittingness. I therefore offer a reminder and a concession. First, remember that fittingness will largely overlap with actual instantiation. My main reason for introducing the notion of fittingness in this investigation is to pick out from all of a thing's properties just that subset which will actually be relevant for aspectual thought, and to include unpossessed properties which may still be relevant. Unpossessed properties will usually be relevant either because they are commonly associated 'mistakes' or because the individual or kind characterized is fictional. In the context of metaphorical communication, this use of fittingness should suffice. I have suggested, but by no means argued, that it will not suffice for explaining our engagement with fiction and other narratives.

If one still simply cannot accept the notions of fittingness and of centrality as I've developed them, even in describing how we often think (as opposed to describing the true metaphysics), then one *could* systematically translate my account in the following way. First, replace all talk about the centrality of features to individuals' actual characters with talk about theoretical centrality in our network of beliefs. Prominence and centrality would then be on a metaphysical par. Drop talk about depth altogether. Second, where possible, replace talk about fittingness (in particular, about the fittingness of affective attitudes) with talk about dispositions to act or to evoke responses in others. As I argued in §3.4, this will not always be possible; in those cases, drop the talk as incoherent.

The ascriptions of patterns of thought that result from this translation will, I think, systematically misrepresent our actual ways of thinking. They will fail to do justice to the full range of our intuitions about characters, and to the objective status that we take those intuitions to have. Further, by eliminating fittingness, the ascriptions will make us look as if we really do think that people and things really have

dispositions which, I maintain, we don't in fact think they have. As a consequence, the modified theory will ascribe to us a willingness to move blithely from the claim that someone or something really does have a certain disposition to the claim that they really do have some other disposition, where there is no good deductive or even probabilistic reason to make this transition, but only an assessment that it would be fitting for those features to go together. These distortions will be particularly pernicious when it comes to the fittingness of affective attitudes toward the subject characterized. It is true, I think, that we sometimes conflate fittingness and fact, or, more often, fittingness and probability. (I return to this point in §3.8.2, where I contrast characterizations with prototypes and stereotypes.) But we often do distinguish them, at least implicitly. In any case it would be useful to have a tool for distinguishing explicitly among these possible commitments. So I think there are serious drawbacks to modifying my account in the manner just suggested. Despite these serious drawbacks, the fact that the revised ascriptions would be kosher from the perspective of a standard theory of mind may be a sufficient benefit for some philosophers.

Finally, one might worry that in ascribing such baroque patterns of thought to us, I've wildly over-intellectualized our actual mental lives. On this score, I think characterizations are no more worrisome than concepts. In both cases, a distinction needs to be made between what is involved in *having* the mental structure and what is involved in *describing* what it is involved in having it. We have and employ concepts and characterizations *in* representing things and kinds in the world; we rarely represent either concepts or characterizations themselves *to* ourselves. The facts which need to be explained in both cases are our dispositions to form certain sorts of judgments and to move from certain thoughts to others. The explanation proceeds by postulating an underlying structure, but that doesn't imply that we must be explicitly aware of that structure itself, or even that the structure must itself have any independent reality above and beyond the relevant collection of dispositions.

3.6.2: *On the Resolution of Disagreement about Character and Coherence of Character*

I've just argued that we often really do undertake the representational commitments that I've built into characterizations: we really do think that a thing's features are structured with respect to each other so that some features are more central than others, and we really do think that some of its features are fitting for it where others would not be. One fact which might make one suspicious of the claim that characterizations do have objective representational import, however, is that disagreements involving them cannot always be brought to a conclusive resolution. As I've described it, the judgment that some feature F fits a's character is typically a judgment that some *second-order* relation holds between F and a. Given the particular sort of higher-order property that fittingness is, judgments of fittingness can systematically come apart from judgments about the lower-level facts without there being any direct incompatibility between those judgments. In this section, I say why such lower- and higher-level judgments can come apart in this way, and why this shouldn't be taken to undermine characterizations' objective import after all.

Lower- and higher-level judgments about character can come apart without any direct incompatibility in two ways. On the one hand, agreement about which lower-order properties a possesses are always compatible with conflicting assessments about F's fittingness to a's character. Both of us might agree that a is in fact F, but I might dismiss F as a fluke while you think it's central, or fits into a's character only at the margin. On the other hand, disagreement about which properties a actually possesses can be compatible with agreement about a's overall character, so long as the disagreement doesn't affect judgments about which features which are thought to be fitting, and about the structure in which those fitting features are located. Because of this systematic gap between higher-order facts about fittingness and lower-level facts about actual possession, disagreement about character can rarely be resolved just by introducing a further lower-level fact. Each interlocutor has the option of dismissing that fact as out of character.

However, there *are* principles for resolving such disagreement, in a way there are not such principles for resolving, say, conflicting assessments of beauty. A characterization that minimizes the number of possessed features which are treated as unfitting, and which maximizes the interconnectedness

of possessed and fitting features, is generally taken to be a better characterization. That is, we allow that some features may be unfitting. But as the evidence increases that *a* possesses many features that count as unfitting on a given characterization, the pressure to revise that characterization so as to maximize coherence also grows.[28] If someone can present a more coherent, more inclusive characterization, then we should accept it.

The criteria we employ in assessing characterizations are thus the criteria of *interpretation* rather than of deductive argument. But this is typical of many areas of inquiry: not just aesthetic analysis, but psychological explanation and scientific theory confirmation more generally. Confirmation occurs at the level of whole theories, not of particular hypotheses. We also have ways of resolving debates about the appropriate structure and fittingness of elements *within* a particular characterization. We offer reasons for thinking that one feature is more central than another, by pointing to other features which would be explained if that feature were taken to be central. We offer reasons for thinking that a given feature is fitting or unfitting, by pointing to other features which fit with it — or by pointing out that there are no or very few such features.[29] And we often succeed in convincing one another through these arguments.

The fact that there are such means for resolving disputes about characterizations, at both the constituent and the overall levels, in turn gives us reason to think that differences in structure and in attributions of fittingness are real differences in how we represent the subjects characterized as being.[30] At the very least, it establishes that characterizations are not merely founded on idiosyncratic causal associations: on which features most strongly *remind me* of which others. In a debate about which

---

[28] This sort of mismatch between fitting and possessed features has traditionally been taken to be less permissible for fictional characters, simply because all of the features they possess have been given to them by an author, and in general authors are expected to attribute only fitting features to their creations. Contemporary authors have intentionally flouted such expectations by endowing their characters with absurd or unfitting characteristics, but the oddness of the resulting characters and narratives supports the claim that characters are usually expected to possess only fitting features.

[29] One way in which an unpossessed property may be included in my characterization of an individual or kind is via my belief that the property is commonly attributed to the subject characterized, even though this attribution is by my lights a 'downright mistake'. How can it be correct to include what is by my own lights a 'downright mistake' in my characterization? It can't. Such 'mistakes' still play a role in metaphorical communication because the claim that the individual does possess that feature may be a conversational presupposition. See §5.1 for further discussion of this point.

[30] Cf. Steve Yablo's discussion of whether modal ascriptions represent objective facts, in "Is Conceivability a Guide to Possibility?"

features are fitting for or are most central to a given subject, we do not accept as evidence the fact that this feature always makes our interlocutor think of another. We require that he provide some reason for thinking that these two features should go together, that they are connected in some independent, intersubjectively available way.

However, it is also true and important that for characterizations the ideal of consistency carries less weight and is less attainable than for scientific theorizing. We do always have the option of dismissing or downplaying certain features as relatively marginal or as out of character altogether, and we do not have this option in scientific inquiry. In science, data is data, and it can't simply be dismissed — although one may hope that troublesome data fail to be replicated, and so that the outlying findings can be attributed to experimental error. By contrast, we do have the option of simply dismissing certain possessed features for characterizations, because characters themselves are not always actually coherent in the way we must assume lower-level facts about the world are. Unfitting events just do happen. When they do, we try to construct a new narrative that incorporates the resulting features, perhaps in a more complex or unusual — even eccentric — manner, but this is not always possible. If the unfitting event or feature is sufficiently important to the individual's identity, we may not be able simply to dismiss it as irrelevant to its character. The individual may be left with a fractured character, one which simply does not hang together. Thus, we can be condemned to live with unfitting, non-coherent characters, and so with non-coherent characterizations, in a way that we can't accept inconsistent beliefs.

Beyond these undesirable but potentially insurmountable failures of consistency, it is also true that we don't merely tolerate, but positively cherish the failure to fit together in an easy, predictable way, as we don't value similar tensions in a structure of lower-level factual beliefs. This is especially evident in our assessments of art objects, but I think it often holds true for people as well. We are intrigued by the possibility that apparently conflicting elements may turn out to fit together in a more complex pattern.[31] However, I suspect that we value multi-dimensionality of character only when we can establish — or at least can wonder at how to establish — a larger coherence within which the various attributes hang

---

[31] Cf. Alexander Nehamas, "An Essay on Beauty and Judgment."

together. We don't value individuals whose properties are simply a collection of randomly and wildly heterogeneous happenstances; and we tend to think that unfitting properties mar the character of the individual who possesses them.

In sum, characterizations are objectively oriented mental representations, consisting of a complex structure of beliefs and affective attitudes. But because characterizations do not primarily consist of beliefs about which properties the individual or kind characterized actually possesses, but rather about what is fitting for it, the standards for appropriate relations between those constituent beliefs and attitudes are different than they are for scientific theorizing. In particular, although we seek overall coherence for characterizations, consistency is neither attainable nor always desirable in the same way as it is in science.

3.7: *Characterizations, Concepts, and Words*

In this section, I locate characterizations with respect to the more familiar philosophical machinery of concepts and of words. My aim here is to demonstrate both that characterizations play a distinctive and needed role in our cognitive lives, and that because they are bound up with our use of concepts and words, that role is a pervasive one. In order to make the contrast clear, and simply to reinforce my claims up to this point, let me briefly recap the most important features of characterizations as I have developed them.

Characterizations are supposed to be a sort of mental representation on which aspectual thought can operate. The main requirements for such a mental representation are, first, that this mental representation be *about* a single subject, but have a structure which could be restructured so as to alter how the subject was thought about without changing what was thought about entirely. Second, its representational content needed to be richer and more nuanced than the content of a single proposition. Third, its content needed to contain the resources necessary for delivering intuitively appropriate metaphorical interpretations.

The primary commitments of my notion of characterizations, designed to meet these three requirements, are as follows. Characterizations are complex representations of a given subject's

character; they consist of structured constellations of beliefs and other attitudes about what is fitting for that subject. The structure varies along the three dimensions of prominence, centrality and depth. A feature's prominence is a function of its intensity, or signal-to-noise ratio, and its diagnosticity, or usefulness for classifying the subject in a given context. A feature's centrality is a function of how fully it is connected to other features in the subject's character, where these connections are in some loose sense causal. A feature's depth is a function of how difficult it would be to eradicate from the subject's character. Only those features which are represented as fitting the subject's character are included within its characterization. Possessed features may not be fitting, either because they are incidental or irrelevant, or because they conflict with central features of that character. Unpossessed features may be included as fitting, because they are generally thought to be possessed and otherwise accord with the characterization, because the subject is fictional or mythical, and so doesn't actually possess any features, or just because the feature in question accords so well with the rest of the subject's character. (This last class of unpossessed-but-fitting features is relevant for fiction, but not so much for metaphor.) Among those features which may be fitting, we find both general and highly specific experiential properties and affective attitudes.

3.7.1: *Intimate Association But No Essential Connection*

The discussion to this point has established that characterizations are conceptual phenomena, where conceptuality is understood in a broad sense. Characterizations are intimately bound up with and dependent upon concepts, in several ways. First, at the most general level, a creature that lacked concepts altogether would also thereby lack the sophisticated representational resources necessary for having characterizations. Indeed, there may be creatures who are able to form some concepts of a basic sort, but who lack the capacity to make the complex connections among thoughts which constitute characterizations. Second, a subject who has no concept for an individual, kind, or property cannot think

about that individual, kind or property at all, and so cannot either have a characterization of it or include it as a constituent within another characterization.[32]

Third, characterizations are specifically associated with particular corresponding concepts. A concept and its associated characterization are 'associated' in that both are representations of the same individual α or property Φ under the same mode of presentation. A thinker could not fail to know that her concept and characterization were of the same individual or property, as she could fail to know this of two necessarily co-extensive concepts. Conversely, different concepts of the same property will be associated with at least marginally different characterizations.

In virtue of their intimate association with concepts, characterizations also have an intimate association with the words expressing those concepts. In particular, if two words express the same concept but differ just in their *coloring* or *connotation*, then the difference will typically lie in the characterizations which are more or less conventionally associated with them. (In fact, I think many purported cases of mere difference in coloring do involve conceptual and even truth-conditional differences after all: in Frege's paradigmatic example,[33] I think 'steed', 'nag', and 'horse' have distinct but largely overlapping extensions. Discourse markers like 'but' seem to be the best candidates for terms that differ only in coloring.) That is, if the two words express the same concept, then by hypothesis they can't differ *either* in what property they ascribe to the subject under discussion *or* in what someone would have to know in order to ascribe that property to the subject. So the remaining difference would seem to lie in the further features which, by using one of the words, the speaker suggests (but merely suggests) that we should expect the subject to have, or in the further affective attitudes she thereby suggests we should have toward it. And this sounds a lot like suggesting a characterization.[34]

---

[32] I include within concepts here what Evans calls Ideas: singular modes of presentation of individuals which have no (or extremely minimal) associated predicational (descriptive) content. There is a more restricted use of 'concept' which applies only to essentially predicative constituents of thought; I employed this more restricted sense in §3.5.1, in the discussion of demonstrative concepts.

[33] "Thoughts," p. 39.

[34] In order to establish the dependence of coloring on characterizations conclusively, I would need to explicitly reconcile my account of characterizations with philosophical theories of coloring. Theories of this sort are not, however, abundant.

The intimate association between words and characterizations also goes beyond just affecting what we suggest by choosing the words we do. Words and the semantic relations between them in turn affect the structure and content of our characterizations. That is, the 'semantic field' — the collection of syntagmatic relations, contrast sets, and affinity sets — in which a word is more or less conventionally embedded[35] often helps both to structure and to suggest elements for inclusion within the associated characterization. Often, the appropriate order of explanation is unclear: is the semantic field structured as it is because it needs to reflect our understanding of the world's structure, or do we structure our understanding of the world in light of the semantic field? The answer may be a bit of both. The current point is just that existing semantic relations can help to suggest and to motivate particular characterization relations. This will become especially important in §5.2, when we consider how hearers generate the appropriate aspect on the basis of the sentence actually uttered.

Despite these intimate connections between characterizations on the one hand and concepts and words on the other, it is important that these connections are merely ones of association and not of constitution or inclusion. Characterizations are conceptual phenomena in a broad sense, but it is essential that we retain a narrower sense of concepts strictly speaking. Perhaps the most basic reason for distinguishing characterizations from concepts is that one can (and often does) have a concept of an individual or a kind without also having the complex, rich, structured pattern of thought that is a characterization. One may have just a few beliefs about it, or defer in one's thought to a more knowledgeable expert. Lacking a characterization can't impugn one's competence with respect to the associated concept.

The second reason for distinguishing concepts from characterizations is that concepts are compositional: they must be capable of combining with multiple other concepts in predictable ways that parallel the grammatical compositional rules of a language to form complete thoughts. By contrast, it makes little sense to think of characterizations as composing into more complex characterizations in such

---

[35] See especially Eva Kittay, *Metaphor*; see also Andrew Goatly, *The Language of Metaphors*, and Patti Nogales, *Metaphorically Speaking*.

systematic, predictable ways. My characterization of Napoleon's mistress isn't a function of my characterization of Napoleon in any straightforward way, as my concept *Napoleon's mistress* is a function of my concept *Napoleon*.

Third, as I've emphasized, characterizations don't always run along same the straight tracks of actually ascribing properties to objects as concepts do. Rather, they deal in how it is fitting for a thing to be. Finally, as I mentioned in §3.2 and will discuss in greater depth in the next section, characterizations are much more idiosyncratic and contextually malleable than concepts. As a result, it doesn't make sense to talk about deference with relation to characterizations in the way it does make sense for concepts. Although we often inherit characterizations from those around us (say, we first hear about pointy-headed intellectuals from our childhood neighbors, or about the stentorian professor from our classmates), characterizations are constellations of beliefs and other attitudes which *we* have.

The relation between characterizations and words is less clear-cut. If we accept that what a word means is just the concept it expresses, as philosophers often do, then characterizations can't be part of word meaning either, by the same line of thought. However, if we are interested in the overall systematic uses and effects of language, then there may be good reason to include connotation or coloring within word meaning more broadly construed; and as we've seen, these notions seem to be closely connected to characterizations. Some associations between words and characterizations may even be fully conventional, so that competence in the words' use requires having the associated characterization; this seems to be true of laudatory and derogatory terms.

Even here, though, I think we should resist tying characterizations too closely to word meaning in general. We have good reasons for denying that names in particular have either descriptive conceptual content or connotations as part of their meanings; but some of our richest characterizations concern the people we know and hear about, and so are associated with the names by which we refer to them.[36]

---

[36] This pair of facts leads Kittay to attribute *ad hoc* semantic fields to names 'in a loose sense' in order to account for the role of names in metaphors while retaining the *semantic* status of her theory of metaphor. She maintains that the names in such cases are not occurring "as true proper names" (note that for her theory to work, this claim must apply to metaphors like 'Einstein is the king of the quantum' as well as to cases like 'George is a real Einstein')

Similarly, we can have characterizations of objects and properties for which we lack standing public terms, which we can only refer to as '*that* guy' or '*this* kind of weather'. Characterizations clearly can't be conventionally associated with these terms, because the terms themselves can refer to very different things in different contexts of utterance.

Theorists about metaphor have often gotten worked up about whether metaphor operates at the 'level of the word' or at the 'level of the thing'.[37] They ask: does metaphorical comprehension rely on interactions between and similarities in word meanings or on interactions between and similarities in how things are? The first option gives rise to a 'substitution' theory of metaphor, on which one word substitutes for another; the second option leads to a 'comparison' theory, on which objects are compared. We can now see that, and why, this debate is ill-founded.

Metaphors are utterances which exploit characterizations in order to alter our ways of thinking about the subject at hand, and thereby to make speech acts about them. Although word meaning does not suffice to determine metaphorical meaning, nonetheless the particular words the speaker chooses to utter make an essential contribution to metaphorical comprehension. This is true both because of the concepts those words express and because of the characterizations associated with them. Characterizations in turn are ways of thinking *about* individuals and kinds. In particular, they are ways of thinking about their characters. Characters, in turn, are structured constellations of fitting properties. I summarize these relationships between characterizations, concepts, words, objects, and properties in Figure 5. Once we have gotten clear on these relations, we can see that on the one hand, metaphor operates fundamentally at the level of thought, with characterizations. But at the same time, characterizations themselves typically take into consideration knowledge both about how *words* are used and about how *things* are.

(*Metaphor*, p. 195, fn. 12; see also p. 175, fn.16). Surely the more parsimonious path would be to acknowledge that although metaphor does importantly involve semantic information, it is fundamentally a pragmatic phenomenon.
[37] See, e.g., Monroe Beardsley, "The Metaphorical Twist," Paul Henle, "Metaphor," Kittay, *Metaphor*.

FIGURE 5 HERE

3.7.2: *The Stability and Publicity of Characterizations*

Thus far in this chapter, I have largely talked about characterizations as if they were temporally stable, interpersonally shared mental representations. I think there is significant truth to this. But we also need to do justice to characterizations' contextual malleability and idiosyncrasy. We do each have our own more or less schematic, default ways of characterizing individuals and kinds around us. But it might be better to say instead that we have default *proclivities* to form certain sorts of characterizations across a variety of contexts. Our own characterizations do often overlap significantly with the characterizations of those around us — of those we might call our confreres. And some ways of characterizing things seem to be basic features of human experience. But even in the cases of greatest commonality, we often also find significant differences. By contrast, concepts are much more stable and public than characterizations.

We might think of the job of concepts, at the most general level, as being to classify things in a useful, relatively systematic fashion, and to allow us to keep track of those classifications and the relations between them in a clear way. By keeping track of those classifications and their implications, we are better equipped to preserve truth in the transitions we make between thoughts. This in turn makes us better equipped to undertake actions that are rational, at least by our own lights. In order to accomplish this job, concepts must classify things consistently across contexts. They must isolate just those features that are salient to the classification, abstracting both from merely accidentally associated features and also from our own merely transient, engaged responses. In addition, because we often exchange information with others and act in conjunction with them, different thinkers' concepts must classify things in largely the same way, and must draw largely the same conclusions from those classifications across a variety of contexts.[38]

By contrast, characterizations have as their job, at the most general level, making sense of our experience of the world in its richness and variety. Making sense of things requires getting an overall

---

[38] I think that sharing *largely* the same system of classifications and implications is sufficient to guarantee the publicity of thought, and that this is all we demand or can get in everyday interactions. The demand for precise, common possession-conditions for concepts which any thinker must satisfy in order to be thinking the relevant thought (cf. Peacocke, *A Study of Concepts)* seems to me to be overly onerous.

view of them, making them hang together in a maximally coherent fashion, and attending most to the things that matter most. To capture the richness of experience, characterizations need to include much more detail than concepts do. They are thus invariably more idiosyncratic, because the details of our experiences often differ. Because very different things can matter in different situations and according to different interests, characterizations need to be much more contextually malleable than concepts are. Because prominence in particular is contextually sensitive, the overall structure of characterizations is somewhat context-sensitive as well. (Recall from §3.3 that one of the two factors determining prominence is diagnosticity, which just is sensitivity to what matters in the situation.) When different sets of features are prominent, different assignments of centrality may also be motivated, by motivating an overall pattern of explanation that makes maximal sense of what are currently the most prominent features. (Of course, which of these assignments of centrality are *accurate* is another, context-independent matter.) Also, recall that moods in particular are highly permeant; thus, my own general mood may permeate my characterization of something, and so highlight particular elements within it.

Thus, I'd like my story about characterizations to be, as much as possible, what Andrew Woodfield calls 'ephemeralist': characterizations are, in general, "freshly made to order", not "pre-cooked, frozen" files in long-term memory.[39] Within any given context, they have all the qualities I have laid out above: features, including experiential properties and affective attitudes, are thought to be fitting for the subject in question and are structured relative to one another according to prominence, centrality, and depth. Which features are thought to be fitting will supervene in some complex way on one's beliefs about which lower-level properties the subject in question possesses. Those beliefs presumably remain relatively stable across contexts, and so introduce some significant stability into characterizations. They thereby explain our default proclivities to structure those features in certain ways. So there will be important cross-contextual continuities in our characterizations. But from context to context, the particular elements included and their precise structure may differ significantly.

---

[39] "Conceptions," p. 565.   See also Douglas Medin and Andrew Ortony, "Psychological essentialism."

Sometimes, our characterizations are 'cemented' by our having actively reflected on the overall character of the subject in question. We might also have settled on just a few central features of a thing as fixing what 'really makes it tick', and leave the other features of our characterization to be relatively context-dependent. As we will see in chapter 4, the application of an aspect can be useful and powerful precisely because it imposes an explicit overall organization on a characterization and thereby provides us with a ready-made way to characterize that subject across contexts. But usually, we just go around with relatively inchoate, schematic, and flexible characterizations, which we fine-tune to fit the situation at hand.

The contextual malleability of my characterizations often stems not just from differences in my own interests and expectations in different situations, but also from my sense of the interests and expectations of others within a conversational context. Different institutional settings, power dynamics, histories, topics, aims, and individual commitments of the parties involved will all affect what I take to be important and appropriate in a particular conversation, and this will affect my characterizations of the relevant subjects in that situation.[40] At a minimum, I need to keep track of how my characterizations match up with those of my interlocutors. If I have evidence that my taking a certain feature F to be particularly central or prominent in a's character is an idiosyncratic feature of my own thinking, and if this idiosyncrasy would be relevant to the conversation, then I need either to make this fact explicit or find some other basis for proceeding. As things go, we seem to be remarkably good at identifying which of our unspoken assumptions are idiosyncratic and which are mutually shared, and relying only on the latter as unspoken assumptions.

These points about conversational malleability should sound familiar. Robert Stalnaker emphasizes the complexity and the shifting nature of pragmatic presuppositions.[41] I could also almost

---

[40] See Goatly, *The Language of Metaphors*, ch. 10, for discussion of the impact of social setting on metaphorical interpretation.

[41] Cf. e.g. "Assertion," "Pragmatics," "Pragmatic Presupposition". Notice in particular that Stalnaker requires, in order for P to be a conversational presupposition, only that the speaker represent herself *as if* she believed that P. She need not actually believe that P, nor need her hearers take her actually to believe it; and they may mutually know that neither believes it ("Pragmatics, pp. 39-40). Weakening the requirement thus goes some (but not all) of

appropriate here Davidson's talk about 'prior' and 'passing theories' of interpretation[42] (indeed, we might think of these 'theories' as themselves characterizations: of the characterizations of our interlocutors).[43] The crucial difference is that unlike Davidson, I think this malleability applies to characterizations but *not* to concepts or to the words that express them. Given what I said above about the basic job of concepts, I think it is appropriate, and important, that we reject the malleability of concepts. In everyday conversation, we do respect the stability of concepts and so of word meaning. We do this by distinguishing between what people strictly speaking *did* say, given the concepts their words conventionally express, and what they tried but *failed* to say (as in a malapropism), or what *else* they meant by what they intentionally said (as in metaphor).

3.8: *Characterizations and Other Mental Machinery*

In this section, I contrast my notion of characterizations with several other, related notions which have been developed with an eye toward psychology rather than primarily toward linguistic comprehension. These are what philosophers sometimes call 'conceptions', and with three sorts of mental machinery more familiar from psychological than from philosophical discussions: prototypes, stereotypes, and scripts.

---

the way toward addressing the departures from actual commitment to the claim that *a* is *F* which have led me to the notion of fittingness. Josef Stern (*Metaphor in Context*, ch. 4) also invokes the notion of presupposition to explain metaphor, but he says nothing more about *which* presuppositions are relevant than that they are those "that must be presupposed to be m-associated with the literal vehicle because, in their absence, the utterance to be determined metaphorically would be metaphorically uninterpretable" (p. 114), where 'm-association' is defined as the "sterotypical, normal, 'metaphor-historical', and exemplified features" that language users know *could* be relevant to metaphorical interpretation. I think it should be counted as a significant advantage of my account that I have developed an independent, unified, and non-*ad-hoc* framework for identifying what sorts of features can be relevant to metaphorical interpretation.

[42] "A Nice Derangement of Epitaphs."

[43] Thanks to Stephen Schiffer and Paul Boghossian for making this point.

3.8.1: *Conceptions*

What I take to be the usual philosophical notion of *conception* is articulated by philosophers like David Wiggins[44] and Jerrold Katz.[45] The core idea is that any given thinker associates various additional beliefs with any given concept, where these additional beliefs are not constitutive of the concept itself and may be relatively idiosyncratic to the thinker. As discussed by Wiggins and Katz, conceptions are concomitants to a quite traditional theory of concepts, on which concepts are definitions comprising the necessary and sufficient conditions for an object's instantiating that concept. Like many philosophers, I would prefer to abandon a sharp analytic-synthetic distinction, and to rank material inferences according to degrees of strength, assigned according to their relative importance for possessing the concept in question.[46] On such a modified view, we could still retain a distinction between concepts and conceptions, but the boundaries between them, and between conceptions and mere factual beliefs, would be fuzzy rather than sharp.

Conceptions seem to be significantly different from characterizations. On a traditional understanding, they have at most a single structural dimension, of theoretical centrality, which isn't equivalent to any of the three dimensions that I've attributed to characterizations. They are composed entirely of beliefs about properties that objects actually or probably instantiate, and so nothing like fittingness is relevant. Finally, affective attitudes have just the same status relative to conceptions as they do to concepts. Different positions are available here, but traditional theories often exclude associated affective attitudes from conceptual content. By contrast, I've insisted that affective attitudes play a pervasive and important role in characterizations.

However, once we give up the traditional view, on which there is a sharp boundary between inferences which are constitutive of the concept itself and those which are merely more or less strongly associated with it, then it is much less clear that conceptions play any distinctive cognitive role or do any real explanatory work. If all 'conception' then means is 'associated beliefs not constitutive of concept

---

[44] *Sameness and Substance.*
[45] *The Philosophy of Language.*
[46] Cf. Churchland, *Scientific Realism and the Plasticity of Mind*, Sperber and Wilson, *Relevance*.

possession', then we do have conceptions, but they're not very interesting psychological constructs. Many though not all of those associated beliefs will also be constituents of characterizations. I have imposed significant additional constraints on what count as the constituents of characterizations and the relations between those constituents, in order to isolate a mental representation which can explain how aspectual thought and metaphorical comprehension work. And I have argued that both characterizations and aspectual thought play a role in our mental lives apart from metaphorical comprehension specifically. If someone can defend the claim that conceptions *per se* do their own explanatory work, then we will need to distinguish between characterizations and conceptions. Until then, I think we should assume that concepts and characterizations together do all the theoretical work that conceptions were originally intended to do.

Andrew Woodfield has sketched a rather different account of what he calls 'conceptions', on which they do have a distinctive role to play. His account is closer to my story about characterizations in several respects than the traditional view of conceptions is. Woodfield emphasizes, as the traditional view does not, that the ordinary use of the words 'conception' and 'conceive' encompasses both truth-directed and "fanciful" ways of thinking. It also encompasses, he thinks, both "conscious episodes" and "long-lasting conditions," and both voluntary and involuntary ways of thinking.[47] What is common across all this variety is that conceptions serve as the "*mental settings* against which the subject does or thinks something": as Woodfield says,

> conceptions are the sorts of things that flavour the contents of conscious states and influence the direction of thinking. When a subject *S* entertains some propositional attitude toward *x* while conceiving *x* in a particular way, the thought-content is conception-laden ...Knowing how *S* conceived things gives [an interpreter] a fuller picture of *S*'s total state of mind.[48]

This notion of a conception as a "mental background" *against which* a propositional attitude is set, whether or not that attitude is truth-directed, is just the sort of idea that I'm interested in. As I mentioned in §3.7.2, I especially appreciate Woodfield's emphasis on 'ephemeralism'. However, he says very little about what conceptions are actually supposed to be; his main concern is to show that

---

[47] "Conceptions," pp. 551-2.
[48] "Conceptions," p. 550, emphasis added.

psychological theories of so-called 'concepts' are unwarranted as theories of either concepts or conceptions.

In particular, Woodfield says almost nothing about how his conceptions are supposed to 'flavour' and 'influence' the contents and processes of thought — even though, as the quote above shows, this is a definitive feature of them. The most he offers here is the hypothesis that one thought may "remind" us of another.[49] This is another manifestation of the common philosophical assumption that the only alternative to the genuinely normative connection between thoughts underwritten by a concept's inferential role is a merely causal, associative connection. Woodfield does have a methodologically sound motivation for offering this hypothesis. He thinks that psychologists have been too quick to postulate new mental entities (which they call 'concepts') for their purportedly scientific explanations of information-processing, where these postulated entities are in fact no more explanatory or evidentially justified than the notions invoked by ordinary folk theory. Thus, he thinks we should try to get along with just the machinery implicit in our existing folk psychology. As will become clear in the next section, I agree with Woodfield about psychological theorizing about concepts *per se*. However, I think I have also shown that in our folk psychology we *are* at least implicitly committed to something like characterizations. I think, further, that I have shown that more than a 'reminding' hypothesis is required for explanatory adequacy, because mere idiosyncratic association, of which reminding is one species, is not sufficient to capture the detailed and nuanced intuitions we have about characters.

Indeed, we can see precisely how Woodfield is led astray by his assumption that 'reminding' will suffice to explain how conceptions 'flavor' and 'influence' thought. He assumes that when an imaginative conception, such as conceiving of a ballerina as a swan, is episodic and voluntary, then our imaginations are *entirely* unconstrained. In these cases, he says, "the agent is not bound by any prior conception of the object: his fancy is free."[50] This makes perfect sense if you think that the only constraints on thought follow from the search after truth. If that constraint is lifted, as it is in imaginative

---

[49] "Conceptions," p. 566.
[50] "Conceptions," p. 564.

conception, then no other constraint should remain. And then 'reminding' would appear to be the natural mechanism for connecting thoughts in a non-truth-oriented context: given the appropriate agent and context, any thought can remind someone of any other.

But this is not in fact how imaginative conception works. In both the perceptual and the non-perceptual cases, there are genuine constraints on how I can think of or conceive a thing *x* even when I'm not aiming to represent how *x* really *is*. I can't think of or see just anything as anything. More specifically, within the context of imagining *x* as *F*, not just anything goes, either. Suppose I am watching a ballerina and imagining her to be a swan, as Woodfield suggests. Even if swans always remind me of the terrible evening I spent by the swan-filled pond in Stratford-on-Avon, that association plays no appropriate role in my imaginative project. Nothing about the ballerina herself legitimates my bringing in that peculiar personal association with swans. Finally, as we will see in the next chapter, aspectual thought and imaginative conceiving *can* play a role in the search after truth, as in the course of scientific investigation. Thus, we need a richer, more detailed account of how 'mental backgrounds' 'flavor' thought, one which makes sense of the normative constraints that are in play both in truth-oriented and in 'free' imaginative contexts. My story about characterizations answers just that need.

3.8.2: *Prototypes and Stereotypes*

If we turn from philosophers' investigations of mental machinery to psychologists', we find rather different assumptions and methods in play, but some of the same fundamental difficulties. In the last 25 years, experimental psychologists have demonstrated that our thinking about things often exemplifies 'prototype effects'.[51] So, for instance, we think of some sorts of objects (e.g., robins) as being more *exemplary* of a superordinate type (bird) than are other sorts of objects which also fall under that type (e.g., penguins or vultures). We think of some properties (e.g., flying and building nests) as

---

[51] See e.g. Eleanor Rosch, "Basic Objects in Natural Categories," Edward Smith and Douglas Medin, *Categories and Concepts*, Sharon Lee Armstrong, Lila Gleitman and Henry Gleitman, "What Some Concepts May Not Be," and Stephen Laurence and Eric Margolis, "Introduction" to *Concepts: Core Readings*.

more *salient* to being of a certain type (bird) than are others properties (e.g., eating insects, singing).[52] And we are quicker to judge, and less likely to make mistakes, on tests of categorization when the objects and properties cited are prototypical for the type in question.

Those experimental psychologists, and cognitive scientists working in their wake, have also attempted to develop theoretical models of thought to account for these effects. In particular, they typically take their evidence to support a prototype theory of concepts — a theory on which concepts just are prototypes. However, the same familiar points about concepts which I recited in §3.7.2 to distinguish concepts from characterizations also demonstrate that concepts cannot essentially involve prototypes, either. First, prototypicality is not determinative of type membership; and ordinary speakers don't take it to be. Even most children know that vultures are birds no less than robins, and that birds don't necessarily fly, build nests, eat insects, or sing.[53] Even the concept of *being an odd number* exhibits prototype effects, but surely all odd numbers are equally odd.[54] But one of the essential tasks of concepts is to classify objects as falling under certain types. If prototypes don't do this, then they can't be concepts.

Second, prototypes aren't compositional in the way that concepts need to be, if they are to account for the rich generativity of thought and for the role concepts play in underwriting inference. To adopt an example of Fodor's, the prototype of a *pet* is, say, a dog; and the prototype of a fish a trout; but the prototype of a pet fish is a goldfish, neither a furry, tail-wagging trout nor a disjunctive dog or trout.[55] Third, prototypes exhibit much the same variability across subjects and contexts that we found with characterizations; but concepts need to be relatively stable across subjects and contexts.[56]

---

[52] Notice that prototypicality is closely linked to diagnosticity, in the sense used by Tversky and appropriated by me as a gloss on prominence. Prototypicality registers the degree to which *a* possesses the properties that are highly diagnostic for being *F*.

[53] See Lawrence Barsalou, "Intraconcept Similarity and Its Implications for Interconcept Similarity."

[54] Armstrong, Gleitman and Gleitman, "What Some Concepts May Not Be," p. 244.

[55] See Jerry Fodor and Ernie LePore, "The Red Herring and the Pet Fish: Why Concepts Can't be Prototypes," and Laurence and Margolis, "Introduction."

[56] See Barsalou, "Intraconcept Similarity and Its Implications for Interconcept Similarity." See also Raymond Gibbs, *The Poetics of Mind*, p. 52ff. for a survey of evidence about contextual dependence of prototypes. Both Barsalou and Gibbs take themselves to be talking about concepts, but this is because they take prototypes to be concepts.

In response to these points, some cognitive scientists (most notably Sharon Lee Armstrong, Lila Gleitman, and Henry Gleitman) have advocated a 'two-factor' theory of concepts. The first factor is what philosophers traditionally think of as a concept: "the systematic categorial description, the sense," which encodes the criteria which must be met for an object to instantiate the concept. The second factor is a more practically-oriented "exemplariness or identification function, used to make quick sorts of things, scenes and events in the world"[57] on the basis of the non-definitive properties that have been most highly correlated with the concept's instantiation in the past. So, for instance, Armstrong et al. say that the categorial description for the concept *grandmother* is "being the mother of a parent"; while the identification function encodes the fact that grandmothers "tend to have grey hair, wrinkles, and a twinkle in their eye."[58] The idea is to have the best of both worlds, preserving the traditional requirements on concepts with the first factor, and accommodating prototype effects with the second.

Notice, though, that this specific analysis is wrong, and wrong in an illuminating way. Perhaps most of us believe that there is a high correlation between being a grandmother and having grey hair and wrinkles — although Clairol and Botox are fast changing that. But it seems highly unlikely that most people believe that the eyes of most mothers of parents twinkle. Neither is it very plausible to claim that we *identify* grandmothers by their twinkles. Of course, Armstrong et al. haven't simply fallen into random error: twinkling eyes are importantly associated with grandmothers. I submit that the association in question, though, is one of *fitting*. Having twinkly eyes is part of our common characterization of a grandmotherly character, perhaps even a quite prominent feature (though not a central one). As a result, it would be a good idea, in filming or drawing an advertisement that involved a grandmother, to make the grandmother's eyes maximally twinkly. The authors wanted to capture this intuition: that, as they say, "some grandmothers seem more grandmotherly than others." But because they were working with just two options — what's *criterial* for instantiating a concept and what's contingently but *typically true* of

---

[57] "What Some Concepts May Not Be," pp. 247-8.
[58] "What Some Concepts May Not Be," p. 248.

instantiations of it — they were misled into locating the twinkly-eyes attribute in the latter, prototype component.

Prototypes seem to be intended to serve two purposes, which are not always coincident: of capturing a rough-and-ready 'identification function' for class membership, and of encoding our intuitions about the characters of certain types. If we think of 'prototype' as a theoretically acceptable euphemism for 'stereotype', as it often seems to be, then the distinction between these two purposes becomes clearer. We all know that stereotypes may not track the genuine probability of some set of features being correlated with the concept's instantiation. Indeed, we often use the word 'stereotype' in everyday speech precisely to suggest that the representation may not be accurate. Nonetheless, we still do *have* stereotypes, and they do influence our intuitive expectations about what we will encounter in novel situations — at times, much to our chagrin.

When the dust settles, I'm not sure what status prototypes will or should have. The so-called 'prototype effects' are real, and we need a model of thought which accommodates them. However, it's much less clear just what these effects are effects *of*. There is a growing consensus, not just in philosophy but also in psychology, that they cannot be accommodated within a theory of concepts *per se*, for the reasons cited above. What, then, about the relation between prototypes and characterizations? Characterizations are quite a bit richer and more nuanced than prototypes as they are typically described, but they also overlap with them in important ways. In discussions of prototypes, psychologists sometimes invoke contextual salience (that is, my prominence), and sometimes invoke essentiality (that is, my centrality), though they do not always clearly distinguish between them. Psychologists also grant that not all of the properties believed to be possessed by a thing are part of our prototype of it, and that some prototypical features may not actually be possessed. Of course, it was the need to accommodate just this disparity between possession and association that led me to develop the notion of fittingness for characterizations.

If we want to distinguish between prototypes, as a technical notion, and stereotypes, as a commonsense one, then these points apply even more clearly to stereotypes. I am thus inclined to treat at

least stereotypes as common, rather coarse characterizations for certain kinds of thing, and so to subsume them within my account. By contrast, if cognitive scientists want to insist that prototypes, as they define them, are indeed cognitive tools for the quick identification of *F*-ness on the basis of properties believed to be highly correlated with being *F*, then this would distinguish prototypes from characterizations. Characterizations are not primarily focused on properties actually believed to be highly correlated. However, as I argued above, it's far from clear that this *is* what prototypes are, if we take them to be whatever it is that produces the empirical effects that psychologists have uncovered under the rubric of 'prototype effects'. In order to settle these issues, we need better evidence about what ordinary people take the intuitions that generate 'prototypes effects' to be intuitions *of*, and in particular, of how these 'effects' interact with their beliefs about how things really are. When we obtain this evidence, then we may find that characterizations can subsume prototypes as well as stereotypes. If so, we would only need to invoke two major types of mental machinery: concepts properly speaking and characterizations.

6.1.3: *Scripts*

In the last 25 years, some psychologists and cognitive scientists have also developed the notion of a mental representation called a *script* or *frame*.[59] This has been the focus especially of work in artificial intelligence. The intention that drives this work is to represent the knowledge that people employ in negotiating everyday situations in a form which would enable a computer to understand stories about those situations. The computer is taken to understand the story if it can answer questions about unstated but intuitively obvious facts in the story. For instance, confronted with the following very short story:

(4) John went to the restaurant. He asked the waitress for coq au vin. He paid the check and left.

the computer should be able to determine that John most likely sat down, and that he ate a meal which satisfied his hunger.

---

[59] See Roger Shank and Robert Abelson, *Scripts, Plans, Goals and Understanding*; Marvin Minsky, "A Framework for Representing Knowledge," David Rumelhart, "Understanding and Summarizing Brief Stories."

The hope is not only to make an intelligent computer, but to learn more about how we normally represent the world to ourselves. The hypothesis is that we represent certain repeatedly enacted, stylized routines as scripts, with subroutines that can be carried out in different ways. The scripts can represent common, public situations in which multiple characters play multiple, more or less mutually acknowledged roles, such as 'going to a restaurant', 'robbing a bank', 'being a jealous spouse' or 'being a Good Samaritan'.[60] They can also represent relatively personal scripts in which a single actor carries out a routinized plan, such as 'romancing', or 'prayer'. The same situation may be represented by different scripts for different actors playing different roles. So, for instance, we will have both a waitress-script and a patron-script for the restaurant situation.

As in the discussion of prototypes, these theorists tend to make unsustainably strong claims — especially, reductive claims — about the role their preferred mechanism plays in understanding across the board. In particular, they seem to assume that scripts encode information as a series of propositions related like a flow chart — although it is not always clear just how to distinguish the theorists' own way of representing scripts from how they claim the scripts themselves are represented in the mind. If this is indeed how they think scripts work, then I am especially suspicious of the claim that the representation of even a massively detailed and complex sequence of propositions about which actions tend to follow which others in which situations could suffice for understanding, let alone for enabling an agent to act appropriately in those situations, as scripts are supposed to.[61] Further, our script-based understanding is intimately intertwined with our non-action-oriented knowledge about how things *are*. A complete theory of the mind should make sense of this, but script-based understanding is usually treated as being distinct from factual understanding.

Despite these worries, I find a theory like Shank and Abelson's scripts useful because that it draws attention to the massive body of knowledge we have about actions, situations and the roles people and things play in them. As such, it helps to dispel the rationalist tendency to focus on objects and

---

[60] *Scripts, Plans, Goals and Understanding*, p. 62.
[61] Cf. Bert Dreyfus's discussion of expert computer systems, in e.g. "From Socrates to Expert Systems: The Limits and Dangers of Calculative Rationality'"; and Searle's discussion of the Background, in e.g. *Intentionality*.

properties in abstraction from situations. Specifically, scripts are representations of ways in which it is *appropriate* to act in a given situation. It is obviously perfectly possible to behave in many alternative ways. And in general there is no law or rule of the sort philosophers are familiar with which underwrites this as *the* appropriate way to act. Nonetheless, we have scripts for how it is *fitting* to behave in fulfilling certain roles and in interacting with certain individuals: in a sports stadium or in an audience with the Queen of England, for instance. Thus, theories of scripts and frames provide us with some resources for investigating and articulating one important type of constituent within characterizations, a constituent which is often ignored by philosophical theories of concepts and of the mind.[62]

3.9: *Conclusion*

In this chapter, I have developed an account of a 'system' of thoughts which exhibits the sort of structure we need in order to make sense of aspectual thought on analogy with aspectual perception. The most important feature of aspectual perception was that thought plays a structuring role in perception itself, and doesn't just accompany perception: aspectual perception crucially involves 'thought in perception' and not just 'looking plus thinking'. The rich, complex structure of characterizations provides us with the resources to explain what the analogous phenomenon of aspectual thought — of 'thought in thought' — might be: we can suppose that one thought is brought *into*, and thereby restructures, a characterization.

I've also tried to ensure that characterizations contain the resources necessary for delivering intuitively appropriate metaphorical interpretations. The notion of fittingness is intended to filter out irrelevant or inappropriate but actually possessed features, and to include in-fact unpossessed but otherwise relevant features. In developing the notion of fittingness, I have argued, contra Black, that

---

[62] One enticing possibility for a theory of scripts is that they might explain pragmatic constraints on the felicitous introduction of definite descriptions without prior reference to individuals satisfying the descriptions. So, for instance, contrast the short story (4) above about a restaurant, with the structurally analogous 'story':

(4') John went to a park. He asked the midget for a mouse. He picked up the box and left.

The definite descriptions in (4) but not (4') are felicitous. Shank and Abelson suggest that scripts "allow for new references to objects within them just as if these objects had been previously mentioned; objects within a script may take 'the' without explicit introduction because the script itself has already implicitly introduced them" (*Scripts, Plans, Goals and Understanding*, p. 41).

characterizations are not merely based on 'associated commonplaces', but on features which are thought to be *appropriate* in some way for the subject characterized. The fact that characterizations typically include highly specific experiential properties and affective attitudes will become especially important when we get to metaphorical interpretation specifically. In particular, these elements provide much of the vividness that makes talk of perception compelling in explaining metaphor, and which lends metaphor its rhetorical force.

However, it would be unsatisfying and unsatisfactory if I generated a totally *ad hoc* sort of mental representation for explaining metaphorical comprehension. I hope my use of examples has motivated the claim that we really do have characterizations, and that we employ them in our everyday thoughts about the world around us, rather than just in the context of metaphor. To summarize the most essential claim once more: one has a characterization of an individual or a kind of individual if and only if one's beliefs and affective attitudes toward that individual hang together in a complex structure which makes just some features fitting. In chapter 4, I explain how an aspect can restructure a characterization.

CHAPTER 4: THINKING UNDER AN ASPECT

We now have a story about characterizations: that is, about a sort of mental representation with the richness and complexity required for aspectual thought. In this chapter, I explain in detail what an aspect is, and how applying an aspect can makes us think of something in a different way without primarily or even necessarily thinking something new or different about it. As Richard Moran says, "There are more ways of changing someone's mind than changing his or her beliefs."[1] In the first instance, aspects 'change our minds' by restructuring our characterization of something. As a consequence of this alteration, aspects can also end up motivating new beliefs and attitudes about that thing.

In chapter 5, I return to the role that aspects play in comprehending metaphors specifically. Right now, though, I want to provide some non-metaphorical examples of aspect-application. As I said at the outset of my investigation of aspectual thought, we employ aspects in thought not merely in creating and comprehending metaphors, but also in engaging with fiction, historical narratives, and aesthetic objects, and also simply in thinking about the people and situations around us more generally. This is not to say that aspects work in just the same way in all these cases. I return in §4.7.4 to some contrasts between the role that aspects play specifically in fiction and in metaphor — although a full treatment of fiction, let alone of other uses of aspectual thought, is obviously beyond the scope of my current task. But I do think that a recognizably unified phenomenon underlies all these different applications, and so that an adequate account of aspectual thought must at least provide plausible resources for making sense of these other cases as well. Right now, attending to some non-metaphorical examples of aspect-application will help to give us a more detailed feel for the relevant phenomenon without posing all the complexity of metaphor. I'll then delve into explaining how these and other cases work.

---

[1] "Seeing and Believing: Metaphor, Image, Force," p. 100.

139

4.1: *Some Examples ~~Revisited~~*

The simplest cases of 'thinking under an aspect' are the stories, true and apocryphal, that we tell about people in our everyday lives. In some cases, hearing such a story simply adds a new entry to the hearer's set of beliefs about that person. For instance, you may tell me that Bill got transferred to a new job at the company, or that Alice has had twins. I may note these facts, perhaps draw some conclusions from them — that Bill will now be rid of his annoying boss, that Alice has entered a new stage of life — and move on. This does not involve aspectual thought in any interesting way, even though the facts in question may be relatively important.

But sometimes, adding that new belief does significantly reconfigure one's characterization of the person in question. Think here of discovering some particularly surprising fact, whether noble or sordid, about a friend's past: that Bob once rescued two children from a burning fire, or that Carol once joined a Satanic cult, for instance. Such discoveries can put the person's other qualities and history 'in a new light': certain features take on a new significance; other features, which previously seemed marginal or incidental, may now seem more typical of their personality; and previously central features may now appear suddenly out of place.

The accusation that Alger Hiss was a Soviet spy seems to have had this effect on many people's interpretations of his testimony before the House Un-American Activities Committee in 1948 and his trials in 1949.[2] Hiss's elegant demeanor, and his tireless work on behalf of Roosevelt's New Deal policies, which had previously seemed so admirable and patriotic, now came to seem like manifestations of elitist snobbery and of dangerous anti-capitalist sympathies. His role in the Yalta conference of 1945 was transformed from just one dignitary honor among a stream of rising accomplishments into a crucial, sinister opportunity to hand over secret documents. His donating an old car to a needy acquaintance was transformed from an incidental act of generosity into a flagrant act of conspiracy. And his family's Woodstock typewriter was transformed from a normal, everyday possession into a tool for executing treason. Here we see the same sort of reconfiguration of multiple facts, both in their individual

140

significance and in their relations to one another. For those who continued to accept Hiss's innocence, by contrast, the prosecution's harping on a few odd and insignificant details seemed all out of proportion to the mass of evidence that testified to his impeccable character.

Given the right circumstances, even the imaginary examples I cited above as cases that wouldn't normally involve a change of aspect might accomplish this same sort of effect: Bill, whom I always thought of as a stick-in-the-mud, now stands revealed as a mover and a shaker; Alice, who I always took to be a giggly, bohemian free spirit, turns out to be strong enough to undergo the momentous and painful experience of childbirth, and now bears the responsibility of shepherding two young lives through the world. I might respond to a story like this by saying that I didn't know the people in question 'had it in them' to undertake that action or undergo that experience. The story need not even be a particularly dramatic to induce the relevant effect. The same sort of reorganization can also be brought on by discovering a 'telling detail' which may not be particularly significant on its own: for instance, that Fred consistently undertips at restaurants, while making a big show of paying for the entire bill.

A rather different example of what I take to be the same basic phenomenon is the parable of King David and the prophet Nathan.[3] The story goes as follows. King David has taken Bathsheba as his mistress; she is the wife of Uriah, a soldier in David's army. When she becomes pregnant by David, David orders Uriah sent to the fiercest site of fighting and left exposed to enemy attack. Uriah is killed; and David, feeling rather self-satisfied, orders Bathsheba to live with him in the palace. The Lord then sends Nathan to rebuke David; and Nathan tells the following story:

> There was a rich man with many flocks of sheep and a poor man with just "one little ewe lamb, which he had bought and nourished up: and it grew together with him, and with his children; it did eat of his own meat, and drank of his own cup, and lay in his bosom, and was unto him as a daughter." One day, the rich man needed to kill a lamb to feed a traveler, and he took the poor man's only lamb rather than of one his own plentiful supply.

Hearing this story, David becomes enraged at the rich man's behavior, at which point Nathan says to David: "Thou art the man." David then repents.

---

[2] This example was suggested to me by John Searle.
[3] 2 Samuel, chapters 11-12. This parable is discussed as a case of aspect-shift by Ted Cohen ("Metaphor, Feeling, and Narrative," pp. 231-242) and in turn by Josef Stern (*Metaphor in Context*, pp. 260-1).

By employing Nathan's parable as an aspect for thinking about his own conduct, David comes to think of his actions in a dramatically different way. However, he acquires no new beliefs — at least no new first-order beliefs — about what he did or about the effects of his actions. He already knew that Uriah loved Bathsheba; that she was Uriah's only wife; that he, David had many wives and riches; that it's bad to take things without compensation; and so on. Rather, what the parable intuitively does is reconfigure the relative importance of these facts, and thereby also make different emotions appropriate. (Perhaps this reconfiguration can itself be represented as a belief after all; I return to this issue in §4.7.1 and §6.2.3.)

Finally, we have a class of cases involving the use of models in scientific theory formation and investigation. Two examples here are using the solar system as a model for thinking about the atom, and using the computer (and derivatively, the hardware-software distinction, data processing and storage, etc.) as a model for thinking about the mind and cognition. In these cases, the model is not being used to make any claims about how the atom or the mind *is*. For this reason, they are importantly different from the communicative use of metaphorical utterances. Rather, in these cases our characterizations of the solar system or the computer are employed as guiding principles for organizing our thinking about the atom and the mind, respectively. As a result, these models can also suggest directions for further scientific investigation.

Despite the great differences among them, all these examples of aspects share the following important features, which we first encountered in drawing out the analogy between perception and metaphorical comprehension in chapter 2. First, their primary effect is non-propositional. They don't primarily serve to introduce a new thought — although they may do that as well — but rather to provide a new way of organizing one's existing thoughts. Second, despite being primarily non-propositional, the aspect's application makes a real difference in thought: the subject is now represented in a very different way. One indication that this is a real representational difference is that we say that some aspects and not others reveal their subjects in 'a true light'. For instance, I think I now understand what really makes Bill or Carol, or Alger Hiss, tick; David now sees his actions for what they are; it's more enlightening and

142

accurate to think of the mind as a computer than as, say, a charioteer driving a team of horses. A second indication that the representational difference is real is that a such new structure in thought makes certain further thoughts and not others make sense. So, I will now make different predictions about Bill's or Carol's future behavior; David will now act differently toward his poorest subjects; scientists will now undertake certain experiments and not others.

The third, important feature of these examples is that they achieve their distinctive effect by bringing one characterization to play a role in another, just as in seeing-as one thought or concept plays a role in perception. The organizing characterization serves as a 'lens' through or by way of which we think about the subject at hand. Wollheim called this phenomenon 'twofoldness' in the case of seeing pictures. In that case, we are simultaneously aware of both the painted surface and the depicted scene; and we see the scene *in* seeing the surface. In this case, of aspectual thought, we are simultaneously aware both of the organizing characterization and of the subject under consideration. In this chapter, I want to explain exactly how these distinctive features of aspectual thought arise.

4.2: *The Basic Story*

I begin with the simplest case, in which one employs an aspect *F* for thinking about an individual *a*. Aspect-application is typically a two-stage process. In the first stage, one attempts to bring one's concept of being *F* within ~~think about the~~ one's *a*-characterization so that *F* is both fitting for and as central as possible to *a*'s character~~er~~.[4] Bringing *F* within one's characterization in this way doesn't simply require entertaining the 'general' thought that *F* ~~Φ~~ can be made to ~~or does~~ fit into one's *a*-characterization in some way or other. It requires actually fitting ~~it~~ *F* into one's *a*-characterization in a particular way. That is, it requires thinking of *F*'s being fitting for *a* in the 'singular' mode discussed in §3.4. One must, for example, configure one's characterization of Carole so that joining a Satanic cult is just the sort of thing she would do. This

involves locating that feature within a nexus of other features: Carol's tendency to follow authority figures, to rebel against mainstream society, to believe in crystal healing, and so on. Because one must fit *F* into one's *a*-characterization in a way that makes *F* maximally central, this will very likely itself produce further realignments in the structure of one's *a*-characterization: other features will now be seen as explained or produced by *F*'s fittingness for *a*.

However, bringing *F* within one's *a*-characterization, even as a maximally central feature, is not yet ~~Something more than s~~sufficient ~~to~~ for apply~~ing~~ *F* as an ~~an~~ aspect ~~successfully~~. (In §4.5.1, I'll explain why it's not necessary, either. For instance, one needn't — and may not be able to — bring *being the solar system* within one's characterization of the atom. But for now I'm focusing on the simplest case.) Recall that for aspectual perception, it wasn't enough just to apply the concept *F* to some element within one's visual field — even if that was a relatively large and important element. That was merely seeing-as, while we were interested in seeing under an aspect. Aspectual perception required employing the concept *F* to impose a holistic organization on the entire visual scene. The same point goes for aspectual thought: thinking of *a* under the aspect *F* requires, insofar as possible, taking one's *F*-characterization to *govern* the structure among all the other elements in one's *a*-characterization. This is the second, and more important, stage of aspect-application. As a bit of terminological regimentation, I'll call the applying or organizing characterization, of being *F*, the 'governing characterization', and the characterization of the subject being thought about, *a*, the 'subject characterization'.

Taking one characterization to govern another involves molding or restructuring one's subject characterization so that it is as much like one's governing characterization as possible. So, for instance, we restructure our characterization of Hiss so that it's as much like our characterization of *being a spy* as possible; or we reconfigure our understanding of Juliet so it's as much like our characterization of *being the sun* as possible. How do we accomplish this restructuring? In a nutshell, like this. Beginning with the most prominent and central features of the governing characterization, we search for matches to those

---

[4] Why does this stage involve one's *concept* of being *F*, and not just being *F* itself? As I said in §3.7.1, characterizations operate on the level of thought, but involve thought *about* individuals and properties. The only

features within the subject characterization. When an appropriate match is found, the matching *a*-feature's prominence and centrality are then raised to a level that maximally approximates that of its matched *F*-feature. How closely the *a*-feature's new structural role approximates the role of its matched *F*-feature is limited by our basic intuitions about what could be fitting for *a*. While an aspect pushes and pulls at our characterization of the subject, it cannot break our prior characterization of it altogether. We continue with this restructuring process for as long as our interest justifies the level of effort required.

4.3: *Aspect-Application: Feature Matching*

　　The process of uncovering appropriate matches is a complex one, calling for ingenuity and imagination. In this section, I lay out the various ways in which features can be matched. Thus far, I have treated all aspects — those employed in the service of contemplation, and of both literal and metaphorical communication — in the same general way. As far as I can tell, all sorts of aspects can also in principle employ the same basic sorts of matches. However, they are differentiated by the sorts of matches they predominately employ. Thus, in this section the distinctive workings of metaphor will begin to emerge more clearly.

　　In the simplest case, both the governing and subject characterizations contain the very same element, *P*, but *P* plays a different structural role, and thus takes on a different significance, in each characterization. For instance, one may think that both Carol and most Satanists enjoy wearing leather, but this feature of Carol's may previously have been neither prominent nor central in one's characterization of her, but simply an incidental fashion choice. Similarly, one may think that both New Deal Democrats like Hiss and Communists revolutionaries like Whittaker Chambers believe in empowering the disenfranchised. And one may acknowledge this as both a prominent and central feature in one's characterization of Hiss. But one may previously have connected that same feature — the belief in empowerment — to importantly different series of further features in each of the two cases. In particular, one may think that where Communists believe that only a totalitarian worker-led government

---

way for me to think about being *F* is by entertaining my concept for that property.

145

will overthrow the hegemony of capitalism, New Deal Democrats believe that education and employment will enable poor workers to wield significant economic and political power within a fundamentally democratic, capitalist society. These further networks of different consequences thus imbue the original feature with markedly different significance in the two cases.

Many matches are not as simple as identity, however. In particular, the more different the two characterizations are, the fewer identical matches there will be. Because of this, many metaphors, especially novel, poetic metaphors, depend on aspects which invoke very few matches of identity; while every (purportedly) true, literal aspect will involve many identical matches. This is, I think, the single most important difference between metaphorical and literal uses of aspects. What remains importantly common across the cases, of course, regardless of which sorts of matches are employed, is the effect which those matches then have: of restructuring the subject characterization so that the matched features play a structural role that maximally approximates the role which the matched feature plays in the governing characterizations.

When identical matches are not available, features in the F-characterization must be matched to merely *similar a*-features. This can happen in three different ways, each increasingly further removed from the simple case of identity. First, two features P and Q may be matched in virtue of their being located at different degrees along a single defining scale. Second, there may be a *structural analogy* between the role that P plays in our F-characterization and the role that Q plays in our a-characterization. And third, P and Q may be *metaphorically* similar to each other. (This sort of match may seem to be circular; I'll argue later in this section that it need not be.) Finally, under certain conditions features can also be *introduced* into the subject characterization from the governing characterization. I'll take up each of these points in turn.

The first type of similarity that can underwrite matches is similarity in some defining scalar respect, like temperature, wealth, or height. In the case of utterances intended non-literally, the governing F-characterization is typically chosen because its relevant constituent feature P occurs toward the extreme end of such a defining scale. As a consequence of its extremity, P is a prominent feature of F's character,

146

and so an immediate candidate for matching. (Recall from §3.3 that intensity is one of the two factors determining prominence.) When P is matched by the less extreme a-feature Q, we raise Q's prominence, often in part by increasing its magnitude on the relevant scale. So, for instance, I might describe the Dean's mansion metaphorically as 'a Taj Mahal', because the Taj Mahal's extreme palatial luxuriousness makes this quality an obvious candidate for being matched by the presumably less extreme luxuriousness of the Dean's residence. Analogous effects will be produced by describing someone as the richest man on Earth — or as an Einstein, an icicle or a mouse.

We might call this sort of feature-matching 'hyperbolic', because if the relevant F-feature were directly matched to an a-feature of the same magnitude, it would produce an implausibly exaggerated magnitude for that a-feature. However, I think this sort of match can also occur in the case of aspects where the governing characterization is believed to apply literally to the subject under discussion. So, for instance, in the minds of some people, once Alger Hiss admitted to being acquainted with a poor, rather degenerate writer named George Crosley (who may or may not have been Whittaker Chambers himself), this acknowledged feature was itself matched to the similar but more extreme feature of fraternizing with a Communist. The result was that this feature of Hiss's — his acquaintance with a poor writer — took on an increased prominence and a more sinister cast than it possessed before. That is, even if one accepted Hiss's own account of the relationship, on which no spying at all was involved, still the very fact that Hiss would be acquainted with this sort of person called his character into question, precisely because it was differentiated only by degree from the more threatening relationship that Chambers described.

The second, less direct sort of similarity is *analogical*. Analogical similarities occur when sets of elements within each characterization, or even sets of relations between elements, stand in an identical higher-order relation. As Dedre Gentner says,

> an analogy is a way of focusing on relational commonalities independently of the objects in which those relations are embedded.... Objects are placed in correspondence by virtue of their like roles in the common relational structure; there does not need to be any resemblance between the target objects and their corresponding base objects.[5]

---

[5] Gentner, "The Mechanisms of Analogical Learning," p. 201.

147

As Gentner uses the term, and as I will use it, analogies *ignore* commonalities between the lower-level, concrete features of the two objects (for me, of the two characterizations). Gentner offers as an example

(1) The wrestler bounced off the ropes like a billiard ball off the wall;

we have here two cases of the very same structural relation, of a hitting-event *causing* a bouncing-event, although the event in each case is different. The shared structure of the two events is made transparent if we represent them thus:

CAUSE [HIT(wrestler$_1$, wrestler$_2$), COLLIDE(wrestler$_2$, ropes)]::
CAUSE [HIT(cueball, billiard ball), COLLIDE(billiard ball, wall)].[6]

Often the relevant higher-order relation will be one of causation, dependence, or implication.[7] When this is so, then the feature that does the causation or implication, on which the other features depend, is thereby more central than the others.

However, not all analogical relations are of this sort. For instance, thinking of the atom under the aspect of the solar system produces the following analogical matches, among others:

ORBIT (planets, sun):: ORBIT (electrons, nucleus)
MUCH LARGER THAN (sun, planets):: MUCH LARGER THAN (nucleus, electrons);

and these matches don't rely on any sort of dependence, causal or inferential. Other analogies often depend upon social roles — as in the cliched match between hunter and lover, both of whom pursue their objects of desire — and specifically upon social hierarchies. So, for instance, Nathan's parable centrally depends on the following analogical matches:

MORE POWERFUL THAN (David, Uriah):: MORE POWERFUL THAN (rich man, poor man)
ONLY BELOVED OF (Bathsheba, Uriah):: ONLY BELOVED OF (ewe, poor man);

GRATUITOUSLY TAKES [MORE POWERFUL THAN (Uriah), ONLY BELOVED OF (Uriah)]::
GRATUITOUSLY TAKES [MORE POWERFUL THAN (poor man), ONLY BELOVED OF (poor man)].

The more specific and distinctive the higher-order relation is, and the more lower-order elements it relates, as in this case, the tighter the analogical similarity between the matched sets of features.

Numerous studies on analogy[8] demonstrate that people perceive analogies only with difficulty, even when the connection is contextually relevant but not explicit. For instance, when asked first to

---

[6] "The Mechanisms of Analogical Learning," p. 209.

solve a relatively easy problem and then immediately to tackle another, harder problem with an analogous solution, subjects often fail to transfer the structure of the first solution to the second.[9] Our memory thus seems to be primarily oriented toward lower-level similarities rather than higher-order structures. (Experts diverge from this pattern; their expertise seems partly to consist in their having developed explicit and well-organized higher-order memory structures for the relevant topic.) However, when people *do* recognize ~~that an~~ the analogy — for instance, if it is suggested to them — then they can both assess it for 'soundness' (for how tightly isomorphic the structural mapping is) and use it effectively to solve the new problem~~s~~. This combination of our difficulty in accessing analogies and their rich cognitive payoff suggests that ~~Thus,~~ one major interest we have in cultivating aspects is helping us to uncover and exploit higher-order structural relations. By bringing two rather different characterizations together, we limit the number of salient concrete similarities between them, and thereby focus our attention on structural similarities.

The third and final sort of similarity between a feature *P* in one's *F*-characterization and a feature *Q* in one's *a*-characterization is *metaphorical* similarity. For instance, if we think of an encyclopedia in terms of a gold mine, two of the prominent features of gold mines which find matches are *being valuable*, and *involve digging around to get the good stuff*. However, the sense in which ~~encyclopedies~~encyclopedias are valuable is more of an intellectual than a financial one; and the sort of 'digging' required is very different in each case.[10] So too, if we think of Sally as a block of ice, the most relevant attribute of each is *being cold*~~. B,~~ but, (at least on the most intuitive reading) the ~~sense~~ sort of coldness ~~i,~~ is quite different for Sally than for ice.[11] In these cases, there is no obvious single attribute which could cash out *the* ~~univocal~~ sense in which both things are valuable or cold, or involve digging.

[7] "The Mechanisms of Analogical Learning," p. 220.

[8] See "The Mechanisms of Analogical Learning"; Keith Holyoak, "The Pragmatics of Analogical Transfer"; Holyoak and Paul Thagard, *Mental Leaps: Analogy in Creative Thought*; and Philip Johnson-Laird, "Analogy and the Exercise of Creativity."

[9] In one example, the easy problem was how to attack a castle without sending all the troops over one bridge and thereby collapsing it, while the harder problem was how to attack a tumor without destroying the surrounding tissue in any one area.

[10] The example and the point are from Andrew Ortony; "Beyond Literal Similarity," p. 174.

Rather, the individual or kind which the governing characterization characterizes has the relevant quality in a literal sense, and the subject of the subject characterization has it in a metaphorical sense. Notice, though, that both *valuable* as applied to ideas and *cold* as applied to personalities are such common, deeply entrenched metaphors that they might be dead ~~ones~~. If so, <u>then</u> both of these examples would rely upon a similarity of *polysemy* rather than metaphor.

Now, the invocation of metaphorical similarity here is likely to seem troubling, in two ways. First, it appears to make my account of metaphor circular: I'm using aspects to explain metaphorical comprehension, but now it turns out that applying an aspect can require grasping a metaphorical thought. However, this poses no inherent problem. It simply introduces a recursive element into the interpretive process.[12] What we must avoid is the following sort of circularity: in order to understand the metaphorical thought that *a* is *F* we need to think of *a* under the aspect of being *F*, but then in order to think of *a* under that aspect we need to understand the very same initial metaphorical thought. We can allow that a salient element *P within* our *F*-characterization is metaphorically similar to an element *Q* in our *a*-characterization, so long as understanding the similarity between *P* and *Q* doesn't in turn require appealing to the similarity between *a* and *F*.

The important requirement is that each embedded metaphorical connection between the constituent features *P* and *Q* must be less complex and easier to grasp than the initial one between *a* and *F*. If so, then we should eventually 'bottom out' into simple, obvious matches. Whether the connection between *P* and *Q* is easily accessible enough to be worth uncovering will depend upon our level of interest in cultivating *F* as an aspect for thinking about *a*. However, the metaphors embedded within aspects will almost always be like those above: very familiar, heavily entrenched, and often widely shared across cultures.

We might well worry whether such highly familiar, entrenched metaphorical matches can themselves be analyzed in non-metaphorical terms. Perhaps these metaphorical similarities simply reflect

---

[11] ~~See~~ <u>The example and the point are from</u> Searle, "Metaphor," p. 87.
[12] Cf. "Beyond Literal Similarity," p. 167.

"brute facts about our sensibility," as Searle says. I do think that such matches can be so familiar and obvious that we would be puzzled if someone asked for an explanation of the connection between them. And when we do try to offer a basis for the match, that basis often appears to be suspiciously artificial and *ad hoc*. For instance, if pressed to cite an underlying basis for the match between personal temperament and temperature, the most plausible candidate would seem to be that cold personalities, like cold things, are relatively inactive and unresponsive. But I'm both less sure whether this holds, and find it less intuitive and informative, than the previous match between cold things and cold personalities. Perhaps the match comes down to a brute intuition that it's *fitting* for people with cold personalities to *be* cold, or to make other people feel cold. This may be right, but, again, it doesn't seem awfully explanatory or reductive.

If there are no plausible subsidiary matches, perhaps this is itself good reason to count the obvious but putatively metaphorical matches as relying on polysemy rather than on metaphor. That is, perhaps there's a literal sense in which personalities are cold, and this is importantly related to but distinct from the temperature range. I myself still think that these matches are at least marginally metaphorical. And I am not particularly bothered by the sort of circularity — or rather, the failure of full reduction — that results from treating just these highly familiar and obvious metaphorical matches as primitive. We would still have made enormous progress in understanding aspects and metaphor if we reduced them to just these 'basic' metaphors. However, if one does find this prospect theoretically troubling, then one can either rely on the (relatively artificial) subsidiary matches, or accept the obvious metaphorical matches as being direct — that is, literal though perhaps polysemous — after all.

The second worry we might have about metaphorical similarity is that it is simply analogical similarity by another name. If so, then my account is just that much simpler. However, the difference between the two, as I see it, is that metaphorical similarities are more multi-faceted and less tightly structured than analogical ones. Analogies *ignore* all similarities among lower-level concrete elements except the structural relations into which they enter; metaphors exploit just such similarities, including especially similarities among highly experiential features and affective attitudes. Analogies depend upon

151

tight isomorphic mappings between sets of elements; metaphors require only that some sufficiently rich set of mappings can be established, whether or not they are mutually consistent. Aspects themselves will thus be more analogical to the extent that they exploit a few overarching abstract analogical structures.

In addition to these four ways in which features in the governing characterization can be matched to features already present in the subject characterization, it is also possible for features may also to be *introduced* into one's *a*-characterization.[13] Not just any feature can be so introduced. First, the introduced feature *P* must be quite prominent in our *F*-characterization; it helps if *P* is one of the *only* prominent *F*-features. Second, *P* (or some suitably de-hyperbolized version of it) must be quite fitting for *a,* on our current characterization of it. Third, some contextual priming — typically but not universally, conversational principles and presuppositions — must lead us to expect that the aspect might introduce a feature in this way. So, for example, someone might utter

(2) Bill's a bulldozer

expecting her hearer to realize that the most prominent feature of bulldozers which could be fitting for a man of Bill's general character is not allowing anything to stand in his way. Understanding this, the hearer might introduce this feature into his characterization of Bill, where nothing like it previously existed.

Affective attitudes seem to be particularly prone to introduction, but only when other matches independently underwrite a connection between the two characterizations. Describing someone as "a damp kitchen sponge," to adopt Moran's example,[14] will make us feel a certain sort of repulsion or at least disdain for him, but only if that person is already characterized as having some other relevant properties: of being bland in conversation and inert in life, for instance. Likewise, once the above analogical matches between David's and the rich man's actions have been established, it becomes quite

---

[13] Cf. Andrew Ortony, 'Beyond Literal Similarity'"Beyond Literal Similarity," p. 173.
[14] "Seeing and Believing," p. 91.

152

natural to transfer the relevant emotions from the one action to the other: of protectiveness toward the beloved object, and of outrage at the more powerful figure's profligacy.

The greater the contextual priming, the looser the connection between the *F*-feature and the introduced *a*-feature can be. That is, the context may make clear that an *a*-feature of a certain sort (say, having a voice of some as-yet-undetermined quality) should be introduced by way of being matched to some *F*-feature or other. If so, when we discover an *F*-feature which can be analogically or metaphorically matched to some particular value of that general sort (say, being like mud or like gravel), the match may be made, and the relevant feature forthwith introduced. Apart from such contextual priming, we would have had no reason to introduce a feature in this way.

4.4: *Application Methods*

With this specification in hand of the means by which the features in a governing characterization can be used to match or introduce features in a subject characterization, the natural next question is how these matches are exploited in order to produce an overall reorganization of the subject characterization. Once again, my central claim is that as matches are made, the matching features in the subject characterization are raised to a structural role that maximally approximates that of their matched feature in the governing characterization. We can now say that a governing characterization will work as an aspect for thinking about a subject to the extent that it meets the following two conditions. First, it must contain prominent and central features that can be matched to features in the subject characterization in one of the above ways, relying on similarities which are themselves sufficiently intuitive and obvious. Second, in applying a new aspect, the effort of uncovering those matches must be exceeded or at least equaled by the effect of the alterations made in the subject characterization.[15]

Unfortunately, I'm not sure that there's much more to say at a general, philosophical level about just how matches are made. (I will, of course, say quite a bit more about how metaphorical meaning specifically is determined in the next chapter.) As I see it, my task in providing a philosophical account

153

of aspectual thought is to remove enough of the mystery from the notion that it can play a genuinely explanatory role in a theory of metaphorical communication, and not simply serve as a placeholder for *whatever* process delivers the answer we intuitively want. In order to do this, the account needs to contain enough detail that it can be tested against our other theoretical intuitions and our actual engaged performances. I have specified the possible sorts of matches between characterizations, and the effects of those matches, in order in order to meet this charge; most theories of metaphorical interpretation, including Searle's and Davidson's, lack this required level of specificity.

In addition, if aspectual thought is to play a role specifically in a theory of metaphorical communication, then there must be sufficient consistency across thinkers so that speakers and hearers can form reasonable mutual expectations about how the other will apply a particular aspect. Otherwise, we will not have explained either how speakers and hearers routinely do converge on common interpretations, nor how some interpretations could be correct and others unjustified. Thus, the general psychological mechanisms I postulate must be relatively intersubjective rather than idiosyncratic.

However, neither of these requirements — for a substantive account of aspectual thought, and specifically for one on which aspects' applications are in fact broadly consistent across thinkers — shows that every facet of aspect application must be determined by a universal rule, let alone one that is discernable *a priori*. Indeed, it is quite unlikely that aspect application can be cashed out algorithmically. If aspects did work by way of our brains exhaustively and systematically 'scanning the feature spaces' of our governing and subject characterizations — as much of the cognitive science literature on similarity assumes — then aspects should at once require much more mental effort and be much more systematically effective than they in fact seem to be.

Rather, as I've stressed, aspectual thought, and the metaphorical interpretation that employs it, is a paradigmatic interpretive activity. This means both that it calls on our general powers of reasoning rather than a modular processing capacity, and that it requires employing imagination and ingenuity rather

---

[15] Cf. Sperber and Wilson's Principle of Relevance, in *Relevance*.

than just deductive computation. It seems likely that different thinkers have different cognitive styles: some may work methodically while others wait for promising matches to leap out at them. And for each of us, some aspects just 'click into place', while others require sustained, explicit effort. Indeed, we do often end up with somewhat different matches, at least in matters of detail. By contrast, we don't usually end up with even minor differences in, say, depth perception, which is processed by a sub-personal algorithm. Thus, I am suspicious of the claim that empirical psychology (let alone philosophical analysis) will eventually uncover the one 'true process' that underwrites aspectual thought.

Of course, the fact that aspect-application often doesn't *seem* to be a systematic search, that it often *seems* to involve imagination and ingenuity, doesn't demonstrate that the brain isn't in fact working methodically in a rapid-fire but fully brute fashion at the sub-personal level. One might believe on general grounds that all mental processes must ultimately be algorithmic. And just as with all interpretive activity, there must be much more going on at the subpersonal level than we are conscious of. What matters, I think, is that we do often end up with roughly the same sorts of matches, and more importantly, that we have a shared sense of which matches are relevant and important in given contexts. I have tried to give an account which makes sense of this.

With these general methodological points on the table, I do now want to make two claims about how we in fact apply aspects. First, I think we have good reason to accept that the bases for matching which I have specified are the ones we do actually seek out and rely upon. I think this partly because these four bases for matching, along with feature introduction, suffice to explain the intuitive effects of particular aspects. (I will demonstrate this further with an extended pair of examples in §4.5.4, and throughout chapter 5.) Further, these are just the sorts of similarity recognized by empirical psychologists like Tversky, Ortony, Johnson-Laird, and Gentner. Finally, when we are pressed to explain or justify an aspect's effects, or when we are stumped in applying an aspect, then we do appeal to the sorts of matches I've described.

Second, I think that other things being equal, the search for matches proceeds in roughly the order in which I presented them. We first seek matches of identity, and then move on to seeking scalar

155

similarities, analogical similarities, and metaphorical similarities; we only introduce elements as a last resort. I think this in part because each of these modes of matching is in turn less direct than its predecessor, involves more potential variables, and so requires more cognitive effort to establish. I also think this because it seems to accord with the deliverances of actual metaphorical interpretation, as we'll see in chapter 5.

However, Dedre Gentner claims that to have empirical evidence demonstrating that this last claim is false. She argues that higher-order structural matches are generally *preferred* to lower-level matches of concrete similarity. For instance, when asked to interpret

(3) Cigarettes are like time bombs,

subjects tended to cite structural relations, such as 'They do their damage after some period of time during which no damage may be evident'. This was true even though, when asked to describe cigarettes and time bombs in isolation from one another, the same subjects tended to list concrete properties and not the higher-order features.[16]

I am not convinced. If there were a *general* preference for analogical matches, as she suggests, then analogies should be much easier for us to uncover and exploit than empirical evidence reveals them to be. We can explain her Gentner's finding in this case partly by pointing to the lack of direct matches between the most prominent concrete features of time bombs and of cigarettes. (Time bombs themselves can look very different: all that really matters is that they have a timer and wires to some explosives.) Because the two domains are relatively different, we have an initial expectation that salient similarities will tend to occur on a higher-order level.   Gentner admits that comparisons such as

(4) The glass tabletop gleamed like water

invite matches between lower-level, concrete salient properties.[17] Thus, when more concrete, direct matches are available, subjects do often uncover them first. ——— However, more empirical investigation

[16] "The Mechanisms of Analogical Learning," p. 222.
[17] "The Mechanisms of Analogical Learning," p. 206.

is required before we can conclusively establish general claims about the preferred order of matching in aspectual thought. The hypotheses I've offered should thus be taken as avenues for such investigation.

4.5: *Complexities of Aspect-Application*

       Now that my basic story about how aspect-application works is on the table, I turn to several

crucial complications, and to a relatively sustained pair of examples. Once again, the basic means by

which I claim we apply the aspect *F* to our characterization of an individual *a* are the following:

- First, one attempts to bring the concept of being *F* within one's characterization of *a* in the 'singular' manner discussed in §3.4: that is, one attempts to assign to *being F* a particular structural role of prominence and centrality within one's *a*-characterization.
- Second, one molds one's subject characterization as fully as possible in the image of one's governing characterization. This occurs through uncovering matches between the two characterizations, and altering the structural role of the matching feature in the subject characterization so that it maximally approximates that of the matched feature in the governing characterization. The search for matches stops when the level of effort required exceeds the effects produced.
- Matches can be made in any of four ways: identity, scalar similarity, analogical similarity, or metaphorical similarity. In appropriate circumstances, features can also be introduced into the subject characterization, either directly or through similarity to a feature in the governing characterization.
- Other things being equal, the search for matches proceeds in the order cited above.
- The primary distinction between aspects employed as a result of comprehending literal utterances and those employed in the service of comprehending metaphorical utterances lies in the proportion of direct matches of identity as opposed to matches of similarity. Aspects resulting from literal comprehension rely mostly on direct matches and secondarily on matches of scalar similarity, while aspects in the service of metaphor typically rely more heavily on metaphorical and analogical matches.

4.5.1: *Varieties of Aspect*

       So far, I have focused on the simple case in which one thinks of an individual under the aspect of

a general kind, where one thinks that it's at least fitting for the individual to be an instance of that kind,

and perhaps that it's true that it is of that kind. The first important sort of departure from this basic case

concerns what sort of characterization is applied to what. Characterizations are in the first instance

representations of the characters of individuals. But recall that characterizations have the same basic

structure and the same basic sorts of elements, regardless of whether they characterize individuals *per se*

or the way individuals of certain kinds tend to be. This latter category also includes the characters of

those who undertake and undergo certain sorts of actions and events, and the characters of those who are

related in certain ways to others. Indeed, at the limit of complexity, we can have characterizations of

entire situations, where these are complex functions of the characterizations of the (kinds of) individuals in those situations.

Now, in the most basic case of aspect-application, we employ a characterization of a kind as an aspect for thinking about an individual: we think of Alger Hiss under the aspect of being a spy, say, or of Juliet under the aspect of being the sun. But we can also apply governing characterizations of kinds to subject characterizations of other kinds, as when we think of butchers under the aspect of surgeons. We can apply governing characterizations of individuals to subject characterizations of other individuals, as when we apply think of Diana as Lady Macbeth, or as Eva Peron. We can even apply governing characterizations of individuals to subject characterizations of kinds, as when we think of the Presidency of the United States under the aspect of Teddy Roosevelt, or of Winston Churchill. Because there is no essential difference in the structure or constituents of characterizations themselves depending on whether they are of individuals or of kinds, these different combinations of characterizations generate no essential differences in how the aspects work. As far as aspect-application itself goes, the distinction between individuals and kinds is a distinction without a difference. Thus, in particular, the process by which we apply our sun-characterization as an aspect for thinking about Juliet is the same basic process as that by which we apply our Nathan's-parable-characterization as an aspect for thinking about David's actions. In both cases, we seek out matches in the subject characterization for prominent and central features in our governing characterization, and raise those matching subject features to greater prominence and centrality.

The second important source of variety among aspects concerns how richly developed and finely structured each of the characterizations are. Dramatically 'lopsided' aspects produce different cognitive effects depending on which side 'outweighs' the other. On the one hand, we can use a very minimal and sketchy characterization to govern a rather rich and complex subject characterization. So, I might think of my familial relations under the aspect of the solar system: one of us is the sun, others are nearer or farther out in orbit. Here it's precisely the sketchy minimality of the governing characterization that's so useful: all the elements in the subject characterization get structured in terms of one clear organizing principle.

159

On the other hand, we can apply a relatively rich, complex characterization as an aspect for thinking about something for which we currently have only a very sketchy characterization, perhaps none at all. So, I might tell you that the man at the next table, whom I but not you know, is the sort of guy who would compete for the record for climbing the most mountains with the least gear in the worst conditions. In a case like this you transfer as much of the elements and structure from the governing characterization over to the subject characterization as possible, perhaps creating the subject characterization itself in the process. Aspects which are lopsided in this direction thus often have primarily substantive, rather than organizational, effects: they introduce new constituents into the subject characterization.

The third, and most important, way in which aspects can depart from the basic case consists in the difference between what I will call 'internal' and 'external' aspects. In the paradigm case of aspect application, we make the concept of being F (that is, the concept associated with the governing characterization F) itself fit in a central way within our characterization of a. This is what happens when we think of Bill as a quarterback, say, or of Carol as a Satanist. Even if we don't believe Bill ever was a quarterback, we still know what sort of character a real quarterback-type is, and we can make sense of Bill's being just that sort of guy; the same goes for Carol, perhaps with a bit more of a stretch.

Sometimes, though, the concept of being F is itself so alien to our characterization of a that we can't bring F itself within it. This is what happens when we try to think of Juliet as the sun. Unlike with Bill and being a quarterback, here we just don't have intuitions about what sort of person you'd have to *be* to be the sun. Other cases fall somewhere in between. My inability to make *being the sun* fitting for Juliet doesn't altogether prevent me from applying the aspect to her, though. The governing characterization can still do its work 'from the outside'. In such a case, I skip the first stage of aspect-application, and focus on the matching process.

Even with external aspects, there are limits to aspect-application: I still can't think of just anything under just any aspect. For example, I will never succeed in thinking of Juliet under the aspect of

160

my social security number.[18] In order for even an external aspect to work, there must still be a sufficiently rich isomorphism between the two characterizations. Just where the limits lie for something's being too alien to fit but not being so alien as to be totally inapplicable differ from thinker to thinker and from context to context. (I return to aspects which fall at the outer boundary of this limit in §4.7.4.)

The aspects employed for comprehending metaphorical utterances are often, even usually external. Indeed, external aspects often produce *better* metaphors. This is true both because they make it more obvious that the utterance is intended metaphorically, and also because they present fewer options for how to apply the aspect, and thereby reduce the hearer's interpretive burden. Given this common correlation, then, we may be tempted to elevate an aspect's externality into a definitive feature of metaphoricity: perhaps all and only metaphorical utterances rely upon external aspects for their interpretation. However, this condition is not sustainable. Insisting on it would be the precise analogue in the realm of thought of fixing semantic deviance as a necessary condition of metaphoricity. And as we saw repeatedly in chapter 1, not all metaphors are semantically deviant. So too here: not all metaphors exploit external aspects.[19] One might, for instance, describe Alice metaphorically as the mother of our group of friends, or Bill as the quarterback of our office. I take it that describing the Dean's house as a Taj Mahal likewise involves an internal aspect. In such cases, the hearer cultivates both of the stages of aspect-application I described above, as opposed to merely the second stage. I return to the distinction between literal and metaphorical utterances that involve aspects in the context of communication specifically in §5.1 and again in §5.3.4.

---

[18] By contrast, I might be able to think of my email address under this aspect: as an identifying code that I should guard closely, because sharing it too widely brings the danger of identity theft, as something I am likely to forget because it is fundamentally arbitrary, and so on.

[19] Conversely, in certain, relatively rare, circumstances a hearer might have to cultivate an external aspect as a means for comprehending a literal utterance. This will occur if her interlocutor has radically different beliefs about the subject under discussion. For instance, in attempting to comprehend an alternative, mythological cosmology, on which the sun and the moon are siblings, a hearer might have to begin by applying the relevant aspects externally, and then try to 'work her way in' to the speaker's literal, direct understanding of those myths.

4.5.2: *Some Potentially Puzzling Features*

In this section, I want to attend to three important features of aspectual thought which initially appear to be in tension. Briefly, these features are first, that the same governing characterization applies differently to different subjects; second, that the same two characterizations produce asymmetrical effects depending on which one is applied to the other; and third, and that the aspect produces reciprocal effects in the governing characterization. These features seem especially puzzling if we analyze aspectual thought as 'thought *plus* thought', or the mere juxtaposition of two thoughts. On such a view, aspect-application would simply consist in entertaining two different subjects at the same time, and then noticing commonalities between them. This is the sort of account offered by Davidson and by William Gass,[20] among others. By contrast, as I have repeatedly reiterated, my account aims to make sense of aspect-application as 'thought *in* thought', on the model of aspectual perception. Because these three concomitant features of aspects have in turn raised difficulties for many theories of metaphor, accounting for them now will stand us in good stead when we turn to metaphor itself.

The key to reconciling these features turns on understanding the sense in which aspectual thought is *directional*. That is, although there is no important difference between characterizations depending on whether they characterize individuals or kinds, nonetheless there is a sense in which aspectual thought is quasi-predicational: one characterization is applied *to* and governs the other.

The first feature is that the same governing characterization will typically produce quite different effects when applied to different subject characterizations. Compare here the different effects of using the sun as an aspect for thinking about Juliet, Louis XIV, Achilles,[21] Anthony,[22] or an atomic bomb; or of thinking of the Presidency under the aspect of Teddy Roosevelt, Jimmy Carter, George W. Bush, or Alfred E. Newman.[23] When the same governing characterization is applied to each different subject,

---

[20] "In Terms of the Toenail: Fiction and the Figures of Life," and other essays in *Fiction and the Figures of Life*.
[21] Cf. *The Iliad*, I, XIX, XX.
[22] Cf. *Anthony and Cleopatra* I.v.55-59 and elsewhere.
[23] Compare here, too, the different effects of thinking of Alger Hiss or Whittaker Chambers under the aspect of being a spy, and thinking of James Bond under the same aspect. Our different affective attitudes toward each class of spy (those who spy on us and those who spy for us) generate rather different matches in the two cases.

different features become relevant. This fact is what made it impossible to maintain a traditional lexical or connotationist semantic theory of metaphor. Both lexical and connotationist theories of metaphor insisted that there must be a fixed set of features, associated with the words themselves, which provided their metaphorical meaning. But as we saw in §1.1, because those theorists had to claim that all those features were part of the meaning of the relevant term in advance of any particular metaphorical application, including all the features that could be relevant led to a massive explosion in ambiguity.

A juxtaposition account of aspects, like Davidson's, is better off here than the traditional semantic accounts were. This is so partly just because it is a theory of aspects rather than of word meaning, and so no threat of ambiguity looms. But further, such an account can explain some of the variety in the effects produced by the same governing characterization by arguing that, because different prominent features in the governing characterization are shared by each of the different subject characterizations, we will in turn end up noticing different features of the subjects themselves.

However, my account can explain this variety in a significantly more accurate way, because it can explain in detail how the very same element in the governing characterization can be matched to very different features in different subject characterizations. This can come about for two reasons. First, the governing feature can be matched in a different way — that is, by identity, by scalar similarity, or by metaphorical or analogical similarity. So, for instance, if *the sun* is employed as an aspect for thinking about an atom bomb, the sun's radiating light and heat is matched by direct identity, or perhaps by scalar similarity, to the light and heat produced by the bomb's explosion. By contrast, when *the sun* is applied as an aspect to a human being, like Juliet, the sun's radiating light and heat is matched analogically.

Second, even remaining within the same type of match for the same feature — say, analogical matches for the sun's radiating light and heat — we still find a great variety in matches. Applied to Juliet, this feature is matched to Juliet's great beauty and goodness; applied to Louis XIV, it is matched to Louis's wisdom and love for his people; if we now apply the same aspect and the same feature to Achilles, it is matched to his towering rage at Patrocles's death; applied to Anthony, it is matched to his

163

handsome, winning ways, which Cleopatra finds so attractive. Because the juxtaposition account merely appeals to noticed similarities, it cannot explain how these very different matches arise.

Finally, the juxtaposition account has no resources for acknowledging that features can be introduced into the subject characterization by applying an aspect. The only resources it allows itself to invoke are noticing commonalities between features which are known to be possessed by each of the things characterized. It thus can make sense only of our being reminded of long-forgotten features, not of our actually discovering new features. But as we'll see in §4.6 and in §5.1, feature introduction is an important part of what makes aspects useful cognitive and communicative tools.

The second potentially puzzling feature I want to discuss is that aspect-application is *asymmetrical*. Compare here the different effects of thinking of surgeons as butchers and of thinking of butchers as surgeons[24]; or of thinking of George W. Bush under the aspect of being the President of the United States (as the sort of guy who fills *that* office) and thinking of the Presidency under the aspect of George W. Bush (as the sort of office that is filled by *that* sort of guy). A juxtaposition account simply cannot account for this fact, because it explains aspects in terms of bringing attention to shared properties, and sharing is essentially a symmetrical relation. By contrast, the asymmetry follows naturally on my account.

On my view, aspect-application begins with the most prominent and central features of the governing characterization, and seeks matches to *them* in the subject characterization. It is no surprise that the two characterizations need not have just the same most-prominent features. Given this, the features we seek matches *for* will almost always be different depending on which of the two characterizations governs the other. Thus, when we think of butchers as surgeons, we seek matches for the highly prominent and central features of surgeons: their delicate touch, their intense training, their concern with cleanliness, the trust we place in them, their care and respect for their patients. When we

---

[24] The example is from Ortony, "Beyond Literal Similarity," p. 166; the idea of analyzing similarity statements asymmetrically derives from Tversky ("Features of Similarity"); Ortony extends this analysis to non-literal similarity statements, of which he takes metaphor to be an instance.

think of surgeons as butchers, by contrast, we seek matches for the highly prominent and central features of butchers: their rough hacking and chopping, their relative lack of training, their style of hauling large slabs of dead meat around and slapping them down on the block.

The third potentially puzzling feature of aspect-application is that in spite of having this directional asymmetry, aspects can often also end up producing reciprocal effects in their governing characterizations. We think of the sun itself differently after applying its characterization as an aspect for thinking about different things. This is one of the basic intuitions behind 'interaction' theories of metaphor. As William Gass puts it,

> We are inclined to think that in metaphors only one term is figurative — 'mouse', not 'Clifford' [in 'Clifford is a mouse'] — but this inclination should be resisted; it is frequently mistaken.... The terms are inspecting one another — they interact — the figure is drawn both ways. Sometimes the metaphor's stress is heavier on one side than another (as I think it is here), but often the emphasis is nearly equal, as if we were seeing mice through Clifford. This can be determined only in context, and of course it would take quite a context to Clifford a mouse.[25]

I think Gass seriously overstates both the frequency and the degree to which the aspect's 'stress' goes from a metaphor's 'topic', or subject characterization, to its 'vehicle', or governing characterization. It is in his interest to play up this reciprocality, because it is just the sort of phenomenon a juxtaposition account would predict. However, if the 'stress' were equal in both directions, as he claims, then there could be no asymmetry, but we've just seen that there is such an asymmetry, and that a juxtaposition theory cannot account for it. Still, Gass is right that there is often *some* reciprocal effect. On my view, this can happen for the simple reason that seeking and finding a match to a feature P draws our attention to that very feature, and thereby raises its prominence relative to the other features in the governing characterization. It can also happen for the more interesting reason that the process of seeking a match for P can enrich our understanding of P itself, and in particular of the connections between it and other features in the governing characterization. A juxtaposition account can explain aspects' reciprocality, and explain it in largely the way I have; the challenge here is to explain all three features simultaneously.

---

[25] "In Terms of the Toenail: Fiction and the Figures of Life," pp. 67-8.

4.5.3: *Sameness of Aspect*

The first of our three features, that the same governing characterization can apply to different subjects in very different ways, raises the question of how we should individuate aspects. If there can be such dramatic differences, when can we say that *the same aspect* is being applied? Apart from mere taxonomic interest, the answer here will play a significant role at several points in chapter 5: in explaining the minimal conditions for successful metaphorical communication, in explaining how relatively uninformed hearers can still interpret metaphors, and in providing the pragmatic reasons which make Stern's purportedly 'semantically ill-formed' metaphors from §1.1.2.2 difficult to interpret.

There are two obvious but extreme, and therefore unhelpful, answers to the individuation question: (1) we have the same aspect only when every F-feature is matched to the same a-feature in the same way and with the same structural role; or (2) we have the same aspect whenever a governing characterization of the same individual or kind is applied, regardless of to what and how it applies. The first answer respects the delicate nuances of aspectual thought, but rules out the possibility of the same aspect being applied by the same person to different things, and even of the same aspect being applied by different people to the same thing, by setting the bar for identity of aspect extremely high. The second answer preserves a tighter connection with concepts and sentences, and so allows more intersubjectivity, but ignores the quite significant differences among aspects that can be generated by the same governing characterization.

We get a better sense of how to answer the individuation question by attending briefly to a more substantive one: what range of similarity and difference do we in fact find among different aspects which apply the same governing characterization to the same or different sorts of thing, (where 'sort of thing' is understood in terms of the basic sortal concepts appealed to by lexical semanticists like Jerrold Katz[26] and by metaphysicians like David Wiggins[27])?

---

[26] Cf. *The Philosophy of Language.*
[27] Cf. *Sameness and Substance.*

It seems to me that because Juliet and, say, an atom bomb are themselves such different sorts of thing, the differences in how one's sun-characterization applies to each are in turn so great that they shouldn't count as generating the same aspect in both cases. One way to support this claim is by appealing to a point analogous to one about the projectability of concepts. Knowing what it means for a car to be fast does help a thinker to know what it means for a road to be fast, but it doesn't help her to know what it is for a knot to be (tied) fast. The failure of projectability in the second, but not the first, case suggests that the second pair represents a case of ambiguity, while the first pair represents at most a case of polysemy. In much the same way, knowing how to apply *the sun* as an aspect for thinking about Juliet does help us to know how to apply *the sun* to thinking about Rosaline, because many of the same features are likely to be matched in much the same way. But it doesn't help us in knowing how to apply *the sun* as an aspect for thinking about an atom bomb. Indeed, it may even hinder us, by raising to prominence features of the sun which are irrelevant for applying the aspect to an atom bomb. Louis XIV represents an intermediate case.

This evidence about aspects' projectability then suggests a criterion for individuating aspects. We might say that we have the same aspect when the same governing characterization is used for thinking about the same basic sort of thing, just because our characterizations of two things of the same basic sort are likely to overlap significantly in their constituents and structure. However, this criterion is still much too weak to capture the relevant projectability phenomenon. The difference between thinking of Juliet as the sun and thinking of, say, Lady Macbeth as the sun is still sufficiently great that one might have a very difficult time proceeding with the one aspect on the basis of the other. But on most standard principles of individuation, *being a woman* will count as being the same basic sort of thing.

We need a more fine-grained criterion of individuation for aspects. The best option, I think, is sameness of an **'essential description'** for the subject, where an 'essential description' is a highly distilled characterization which specifies the key element that underwrites the thinker's characterization of the subject. Examples of such essential descriptions might be 'beloved girl' (as Romeo thinks of Juliet and once thought of Rosaline), 'ambitious and conniving woman' (as Shakespeare intends for us to think

167

of Lady Macbeth), or 'powerful political figure' (as we think of Louis XIV). Given the same essential description for two governing characterizations, we are virtually guaranteed a sufficient overlap between the two characterizations more generally to underwrite projectability. More generally, in order to count as applying the same aspect, two thinkers must be employing the same essential descriptions of the same things for both their governing and subject characterizations. If they do so, their overall patterns of thought will be sufficiently similar to make communication, at least in broad outline, possible.

Notice here that we may only be in a position to fix the appropriate essential description of the subject *after* we apply the aspect, because the process of applying the aspect may itself alter what we take to be most central to the subject's character. However, when we have a purpose for entertaining an aspect, as in a conversational context, then that purpose will typically already isolate just a few of the subject's features as the most prominent ones for that purpose.

4.5.4: *A Pair of Examples: Juliet and Louis XIV*

In this section, I want to support the claim that a single governing characterization can generate rather different aspects when used to think about different things, by describing in some detail how *the sun* applies as an aspect to Juliet and to Louis XIV. This will also serve as a test case for my account's ability to do justice to rich and fruitful aspects. Because almost any utterance of the corresponding sentences will be metaphorical, the discussion will anticipate and lay some of the groundwork for the treatment in chapter 5 of the content of the corresponding metaphorical utterances. However, it is important to remember that in this chapter, my task is just to explain aspectual thought itself, not metaphorical communication.

I start by describing the outlines of what I take to be our characterization of the sun. I will try to do this in a relatively neutral way, so as not to stack the deck by inserting those features needed to underwrite the two aspects. However, remember as well that in a particular context, our characterization of the sun will likely already be molded in a way that makes it amenable for application to the relevant subject.

168

Here, then, are some of the more central and prominent features of our characterization of the sun: being an enormous fireball; being the center of the solar system; being our primary source of heat and light; enabling plants to grow; disappearing at night and returning in the morning to begin the day; blinding us at least temporarily (in a way that has a distinctive phenomenology) with its brightness when we try to look directly at it or when we come into its presence from a dark interior; making our skin feel warm, and potentially burning it (again, both of which have distinctive experiential components). Less central features might include being millions of miles away and millions of years old. Our characterization likely also includes images of the way the sun looks in different sorts of weather and at different times of day (on a cold grey winter's morning, at the height of summer noon, hanging low on the horizon just before an autumnal sunset, etc.), and its its being fitting for the sun to haveing different moods (of gentle caressing, of blazing parching anger, etc.) and for us to have different emotional responses to it (of gratitude, of fear and seeking protection from it, etc.).

To get our intuitions about Juliet going, recall the extended metaphorical passage from which the famous predication is taken:

But soft, what light through yonder window breaks?
It is the east and Juliet is the sun!
Arise fair sun and kill the envious moon
Who is already sick and pale with grief
That thou, her maid, art far more fair than she.
Be not her maid since she is envious,
Her vestal livery is but sick and green
And none but fools do wear it. Cast it off.
It is my lady, O it is my love!
O that she knew she were!
She speaks, yet she says nothing; what of that?
Her eye discourses, I will answer it.
I am too bold, 'tis not to me she speaks.
Two of the fairest stars in all the heaven,
Having some business, [do] entreat her eyes
To twinkle in their spheres till they return.
What if her eyes were there, they in her head?
The brightness of her cheek would shame those stars,
As daylight doth a lamp; her [eyes] in heaven
Would through the airy region stream so bright
That birds would sing and think it were not night.[28]

---

[28] *Romeo and Juliet*, II.ii.1ff.

Here also for guidance is Stanley Cavell's famous stab at a paraphrase of the core metaphor (neither he nor I try to do justice to all the complexities of this very complex passage):

> Now suppose I am asked what someone means who says 'Juliet is the sun'….I may say something like: Romeo means that Juliet is the warmth of his world; that his day begins with her; that only in her nourishment can he grow. And his declaration suggests that the moon, which other lovers use as emblems of their love, is merely her reflected light, and dead in comparison; and so on.[29]

Because being the sun is such an alien feature for a girl to have, the aspect is external: we can't bring *being the sun* itself within our characterization of Juliet as fitting for and central to it. Because a sun and a girl are quite different sorts of thing, we should expect fewer direct matchings of concrete properties and more analogical and metaphorical matchings.

~~Thus,~~ When our sun-characterization is applied to our characterization of Juliet, its enormousness, its being a ball of gas, and its age find nothing like appropriate matches~~, i.~~ These features must be ignored or else we'll be prevented from applying the aspect fruitfully. Indeed, it seems to be precisely the need to discard features like these, which are so central to our characterization of the sun, that prevents us from bringing our concept of *being the sun* itself within our characterization of Juliet.

The sun's being by far the brightest celestial body in the sky finds an analogical match with Juliet's being the most beautiful girl in all of Verona. The fact that the sun is a natural and original source of light (unlike either a lamp or the moon, both of which figure as contrasting subjects) is in turn directly matched to Juliet's beauty being natural and original. Romeo (or rather, Shakespeare) emphasizes the match between brightness and beauty later in the passage by treating it as not merely one of analogy but of scalar similarity, by comparing Juliet's brightness to those of the stars and a lamp. He then 'super-hyperbolizes' the match between the sun's brightness and Juliet's beauty by insisting that Juliet's eyes and cheeks are not just *as* bright as the stars (already a hyperbolic claim), but far brighter. The initial analogy between the sun's superlative brightness and Juliet's superlative beauty has as a direct consequence a match between sun's making the other stars look pale and Juliet's making the other girls look dull and ugly.

---

[29] "Aesthetic Problems of Modern Philosophy," pp. 78-9.

The sun's being our primary source of light and heat finds an analogical match in Juliet's being an exemplar of virtue. (This may in turn depend on the entrenched metaphorical similarity between light and virtue; if so, the match is metaphorical.) The sun's enabling plants to grow finds a match in Juliet's goodness encouraging Romeo to mature beyond the rather frivolous teenager he has been. I think this is a match of metaphorical similarity: the sun and Juliet both encourage something to grow, though what counts as 'growing' is metaphorical in the latter case. However, one might well take it to be a match of identity, if one thought that the latter case of 'growing' were literal.

The sun's cycle of rising and setting finds an analogical match in Romeo's thoughts beginning and ending with Juliet. The night might in turn be analogically matched to the long periods when Romeo is excluded from Juliet's presence. And the warmth the sun makes us feel (with its particular experiential aspect) is matched by the warmth that Romeo feels on those occasions when he is near Juliet. ('You know how the sun makes you feel on a warm spring day?', Romeo might ask. 'Being around Juliet makes me feel precisely like *that*.') This 'warmth' is both literal (his physical excitement) and metaphorical (his emotional state of happiness).

So much for the aspect of the sun as applied to Juliet. I've summarized these characterizations and the matches between them in Figure 6.

FIGURE 6 HERE

When the same characterization of the sun is now applied to our characterization of King Louis XIV, rather different features come to the fore. We might take as our guide here the following paraphrase of the metaphor of Louis as the Sun King, constructed on the model of Cavell's paraphrase above:

> Now if I were asked what Louis meant when he called himself "The Sun King," I might say something like: Louis meant that his glory nourished the growth of his court and his country; that his person was the pinnacle, if not the center, of the French universe; that his power was greater and more dazzling than any lesser luminary's in the nobility. His self-characterization also suggests that the king's subjects by nature must suffer a certain distance from their monarch, as if their lives were lived in the mundane while his blazed in the firmament, and as if they would be burned if they attempted impertinently to approach too near to his brilliance, and so on.[30]

Now the sun's distance from us, its age, and even its enormousness — elements that were ~~inert~~ ignored in the aspect's application to Juliet — do find matches. The sun's distance from earth is matched, analogically and then hyperbolically, to Louis's distance in importance and hence in accessibility from his subjects. The sun's age is ~~rmatched~~matched, not by Louis's own age, but by the ancient ancestry of his Kingly authority. The fact that the sun occupies a natural place in the heavens is matched analogically to the King's occupying a natural role in the national order. The sun's location at the center of the solar system *might* be matched analogically by Louis's location at the center of the court's, and nation's, life. However, a heliocentric understanding of the planets was just beginning to circulate at the time of Louis' reign, and so it seems unlikely that this characterization was fully established. In a pre-heliocentric view, the sun would still be the most important, biggest, brightest celestial body, although not literally the central one. These features of the sun would then be matched to Louis' importance for the court's and nation's life.

Now the sun's brightness gets matched to Louis's gloriousness, not to beauty. The heat and light that the sun radiates are matched metaphorically by the king's great wisdom and his love for his nation. This match in turn motivates an analogical match between the sun's causing the plants to grow and Louis's cultivating enlightenment in his subjects. The sun's capacity to blind those who come into daylight from a dark interior is matched by the ~~the~~ King's ability to dazzle those not accustomed to the

---

[30] I owe this able paraphrase to Genevieve Gunther.

court. (I suspect these last three matches were the most prominent and central ones on Louis's own understanding of the metaphor.)

Louis would likely have been willing to accept that, like the sun, he could burn those who got in his way. But insofar as he sought to cultivate a gracious and enlightening self-image rather than a threatening one, he would want to downplay this feature, and focus instead on the feeling of gentle warmth we feel in the sun's glow; this is analogically matched to the admiration and security the King's subjects should feel under his benign protection. The sun's disappearing at night and returning in the morning to begin the day must now be ignored altogether: the King is supposed to be a constant beacon for his people, not a mere sometime presence.

I've summarized the two characterizations and matches between them in Figure 7. If we compare Figures 6 and 7, we can see that much of the structure of the characterization of the sun is preserved and transferred over to the characterizations of the two subjects when they are thought of under the sun's aspect. However, we also see that many of the same sun-features find matches to different sorts of features in Juliet and Louis XIV. The analogical and metaphorical matches are underwritten by different relations, or are underwritten by the same abstract relation which is then instantiated by different features in the subject. Finally, we can see that different sun-features must be downplayed or even ignored altogether in the two aspects. This reflects the fact that the two 'essential descriptions' are quite different: a powerful male political figure as opposed to a young female love-object. On the other hand, Louis and Juliet are still both people, and so there is some significant overlap. An application of our sun-characterization as an aspect to something other than a person — say, to an atom bomb— would differ even more in which features were matched and how.

FIGURE 7 HERE

4.6: *Benefits and Dangers of Aspects*

In §§4.2 through 4.4, I laid out the basic ways in which I take the application of an aspect to work, and used that explanation to say when an aspect will work for thinking about a given subject. An aspect works by taking the most prominent features in the governing characterization and seeking matches to them in the subject characterization; matches can be direct or in virtue of scalar, analogical, or metaphorical similarity. Finally, in the right circumstances the governing characterization can introduce a feature into the subject characterization. When matches are found, the structural role — of prominence, centrality, and perhaps depth — which the matched feature plays in the subject characterization is altered so as to maximally approximate that of its matching feature in the governing characterization. An aspect works just in case a sufficiently large number of appropriate matches can be uncovered with a sufficiently minimal amount of effort. In §4.5, I considered some important complexities in how aspects can work and a pair of examples of how they work. I now want to step back a bit, and ask what work aspects do *for us*. That is, why should we be so interested in thinking of things under aspects? I think that aspects have two primary sorts of benefits. The advantages they confer can be considerable, but these benefits are also intimately intertwined with parallel dangers.

4.6.1: *Organizational Efficiency*

The first and most important benefit of aspects is that they produce much more highly structured and coherently organized characterizations of the subjects we think about. Much of what we know about a subject — sometimes seemingly everything — 'falls into place'. As I've stressed, the benefit of imposing an aspect is not typically manifested in explicit thoughts about the structure of *a*'s character, though it can be so manifested. Rather, we *use* that structure in thinking more efficiently about the object itself and about what is fitting for, and likely to be true of, it. This increased understanding is most clearly evident in the way we may come to 'know our way about' in our dealings with that thing: by a new ease and felicity of interaction with it. Because the aspect gives us a principle for assimilating new

176

facts within the framework of our existing characterization, we now have an increased ability to make sense of its behavior in new contexts.

In particular, because some things really do matter more relative to certain contexts and certain interests, an aspect which makes just those features prominent in that context will be more appropriate and useful for thinking about that subject in that context. And because some features really are more central to a subject's character than others — because some central features really do make the subject 'tick' — an aspect which makes the genuinely central features prominent will represent the subject's character in a more accurate way. As a result, it will enable us to make more accurate predictions about how that individual will in fact behave.

To spell out this talk about more coherent and structured characterizations, we might distinguish among four degrees of structure that can obtain in our thinking about some thing or kind. First, and most minimally, one might simply have the ability to pick that thing out in thought, by means of some mode of presentation or concept, but lack much in the way of beliefs or other attitudes toward it.[31] Second, one might associate a mere hodge-podge, or an unstructured list, of beliefs, experiences and attitudes with it. Third, one might have a rough-and-ready characterization of it, where that characterization is not particularly nuanced or coherent, but just workable for present purposes; this is perhaps the typical way in which we think about things. Finally, one's characterization of it might be quite structured and interconnected. Applying an aspect is one of the most effective and convenient ways to impose such a structure.

Often we are driven to achieve a more structured and interconnected understanding as a result of some conflict in our rough-and-ready characterization. Some of the subject's features are apparently in tension with each other, though each seems true and fitting on its own; we thus seek a new structure in which that tension can be reconciled. While it is possible to impose a new, overall organization on a characterization without employing an independent characterization of some other thing, a characterization associated with a standing concept is useful because it brings along its own ready-made

177

structure. It can be useful even if this structure is itself relatively rough, because it still provides a unifying organizational principle for thinking about the subject.

The danger of aspects resides precisely in the fact that they employ a single principle for organizing what we take to be fitting for a particular subject. Most people and things are not in fact neatly governed in which properties they possess by a single overarching concept. As a result, an aspect will almost inevitably ignore or suppress some of a thing's features, and those features may themselves be central to its character.[32] The difficulty is that we cannot tell from our own epistemic situation what is being left out. This is true even when aspects are not involved, but an aspect compounds the danger by requiring not just a maximally coherent characterization, but a coherent characterization organized in terms of a single principle. Many aspects — those intended to serve both literal and metaphorical comprehension — are useful precisely because they show up a few features of the subject's character very clearly. They thereby eliminate what, for present purposes, may count as unnecessary and bothersome noise. But an aspect's very usefulness may then blind us to its inappropriateness, lulling us into a false sense of security about just how robust our understanding of the subject is. We end up focusing on just those few highlighted features, and neglect those that don't fit the aspect. So, for instance, thinking of Alger Hiss under the aspect of a spy raised enough features to prominence to convince many people that he must in fact be a spy. But it also downplayed a significant number of other features, which might still have been important constituents of his character — regardless of whether he was in fact a spy or not. Similarly, I believe that the computer model of the mind — again, whether or not it truly applies literally — leads cognitive scientists to overplay those of our thoughts which fit the deductive computational model, and to neglect the intimate role that action and imagination also play in our mental lives. (I return to this point again in §6.3.)

---

[31] This corresponds to 'thinking of something under an aspect' in the weak, Fregean sense of the term 'aspect'.
[32] Artifacts — and especially simple designs like the duck-rabbit figure — are a notable class of potential exceptions here, because they are often designed with a single concept in mind. Indeed, this partly explains why they illustrate aspect-perception so well (and switches of aspect, when the single concept intentionally involves an ambiguity). Even here, though, it's important to remember that some artifacts, most notably artworks, have been designed to resist full subsumption under a single concept: they have been designed precisely to be complex and multi-valued.

4.6.2: *Feature Introduction*

      The second important way in which aspects work for us is by helping us to uncover new features in a familiar object. This can occur in at least two ways. The first case is not really one of discovering *new* features, but of bringing attention to neglected features. David already knew that Uriah loved Bathsheba, but it was in his self-interest to neglect this fact. The aspect induced by Nathan's parable brought it forcefully to his attention. Thinking of one thing under multiple successive aspects may be especially useful in this regard, because each aspect raises different features to prominence in turn.

      But as we've seen, the aspect may also introduce a new feature into, or at least suggest a new feature for, our subject characterization. This can happen directly, because a highly prominent feature in the governing characterization, P, lacks a match in the subject characterization but still seems fitting for the subject. So, for instance, because planets rotate the sun in fixed orbits, scientists were led to wonder whether something analogous — the electron shell — might be true of the atom. Similarly, thinking of the mind as a computer encourages the search for a language of thought, physically realized in the brain, as an analogue to the computer's hard-wired operational language of on-off bits. New aspects often encourage scientific progress in this way, by suggesting potentially fruitful avenues for investigation.[33]

      Feature introduction can also occur in a more indirect fashion, by reconfiguring our subject characterization so that an entirely new feature, Q, which is not itself matched to any feature in the governing characterization, now seems fitting for it. In particular, the aspect may lead us to notice a property for the first time altogether, and not just the fact that this particular subject instantiates a familiar property. So, for instance, we might pick out a certain sort of gesture for the first time when we think of the dancer we're watching as a swan; but we might then come to notice that same gesture in a host of other situations, by a host of different people. Similarly, we might only distinguish a certain thick, musky taste in a glass of wine when we think of the wine as containing leather, or a certain tart acrid taste when

---

[33] Cf. Richard Boyd, "Metaphor and Theory Change: What is 'Metaphor' a Metaphor For?"

we think of it as containing iron; but we might then develop a reliable disposition to discern those particular flavor in wines generally. The particular gesture or flavor in question may not correspond to any particular feature in our characterizations of swans or of leather or iron, but it may still be brought to the fore under that aspect.

In such a case, the discovery of the new feature is not a direct consequence of the governing aspect, but the aspect guides our investigation of the subject, and so our discovery. I think this 'guiding' role that aspects play in feature introduction, and indeed in thought more generally, is at least nicely analogous to, and in the right circumstances might actually be a case of, what Kant described as "the free employment of the imagination" in contemplating an aesthetic idea: in such employment, he says, thought is neither fully bound by rational law, nor is it simply jerked about by the empirical "law of association."[34] Rather, the imagination here seeks out the most appropriate development of its topic.

Some people also think that an aspect may help us to discover new features in a third way: by providing us with a perspective which is necessary for perceiving that feature at all. They argue that there are features to which we can't in principle have access except by way of an aspect; for this reason, the aspect itself is sometimes said to *create* the property. The idea is that this still counts as genuine access to the way things are because, given this perspective, there can be real agreement and disagreement about whether the relevant feature genuinely applies.[35]

A class of such putative properties is provided by what Wollheim calls 'projective properties'.[36] Their relevant features are first, that "we identify them through experiences that we have, experiences that are both caused by those properties and of them,"[37] and second, that these experiences are themselves *complex*, in that they partly depend upon our having projected, in a quasi-Freudian sense, an intolerable

---

[34] *The Critique of Judgment*, p. 176.
[35] Carl Hausman ("Metaphors, Referents, and Individuality") articulates something like this view. Notice that the radical claim is different from (and considerably more extreme than) Black's claim that "it would be more illuminating... to say that the metaphor creates similarities than to say that it formulates some similarity antecedently existing one" ("More about Metaphor," p. 37). Here it is the new *similarity* that obtains between two previously and independently existing features. Hausman seems to conflate these claims.
[36] "Correspondence, Projective Properties, and Expression"; see also McDowell, "Aesthetic Value, Objectivity, and the Fabric of the World," and "Values and Secondary Qualities"; and Simon Blackburn, *Spreading the Word*.
[37] Wollheim, "Correspondence, Projective Properties, and Expression," p. 149.

mental condition of our own (e.g., melancholy, fear of loss of love, overabundance of joy) onto a part of our natural environment. Having projected our own mental condition onto something around us in this way, we recognize various other parts of nature as being such that we "might have, or could have, projected this or that kind of feeling" onto them as well.[38] When we perceive some part of nature in this way, it "is felt to be, not actually melancholy, but *of a piece with* [our own] melancholy."[39] Thus, the claim is that it is only through having this sort of intolerable emotion, and knowing what it's like to have projected that emotion onto one's environment, that one can be in a position to recognize the particular way that parts of the world must be if they are to underwrite other instances of such projection.

I'm not sure myself what to think about whether aspects can have an essential role to play in the perception of otherwise inaccessible properties. I'm not sure that Wollheim's projective properties — or the simpler projective properties that people like McDowell and Blackburn discuss, such as being good or being funny — should be counted as real properties, though I'm more or less inclined to think they are. I'm much less sure that they really are accessible only via the current cultivation of an aspect.[40] It would help considerably to have an example of a property which clearly met both these requirements, but it seems to be in the nature of the case that proffered examples will not uncontroversially meet both requirements at once.

Again, as with aspects' organizational power, the ability of aspects to introduce features into the subject characterization poses a danger as well as a benefit. An aspect may make a feature *P* seem fitting for *a*, and thereby lead us confidently to predict that *a must* (or cannot) actually possess *P*. That confidence may prevent us from investigating further, or it may lead us to distort information that does come our way. Thus, the solar model eventually came to hinder the investigation of the atom at least as much as it aided it. This danger becomes especially pressing when we conflate fittingness and truth. What appears to be out of character or unfitting for a subject when it is thought of under a certain aspect

---

[38] "Correspondence, Projective Properties, and Expression," p. 153.
[39] "Correspondence, Projective Properties, and Expression," p. 151.

— and indeed, what actually *is* out of character for it — may nevertheless be quite *possible*, even actual. We obviously cannot leap directly from a feature's fittingness for an object to its actual instantiation by it.

Finally, aspects are both powerful and dangerous because of the important role that affective attitudes play in them. This is part of why, for instance, Nathan's parable alters David's understanding in such a powerful and dramatic way, and why parables in general are such effective pedagogical tools. Because attitudes are permeant, they color other features within our characterization and raise the prominence of features that fit with them. Those further features can then seem to confirm the attitude's appropriateness, and thereby make the overall aspect more difficult to dislodge. We can end up with a sort of interpretive feedback loop, which may prevent us from getting a clear view of the facts at hand. For instance, my characterization of my interlocutor's way of standing near me will partly depend on whether I *already* find him menacing, or flirtatious, or simply think he's hard of hearing. But if I do find him menacing, his stance may look to me to be *independent* evidence for his menacingness, even though it might not itself appear menacing if I didn't already think of him that way.

All the dangers of aspect application are evident when we think about the much-maligned effects of stereotypes, which, as I said, can be thought of as coarse, simple, and publicly shared characterizations (see §3.8.2). When we operate with stereotypes, we can leap to unwarranted conclusions about what further features the individual in question must in fact possess; and we can let our affective attitudes toward the type characterized lead our judgments about the facts in a particular case. The unreliability of eyewitness testimony provides evidence, I think, of these dangers. But it is also important to remember that even stereotypes can bring important benefits. At a minimum, we gain an increased base of beliefs about individuals of that type on which we can draw in interacting with them; and we are ready for quick action and response as we are not when we interact with entirely novel types. Careful cognitive bookkeeping — clear thinking *about* our patterns of thought — can ameliorate many of the dangers of

---

[40] It's also far from clear that any of these three theorists think that aspects are required. Rather, they think that these properties are projective in that they require a certain sort of mental capacity or disposition to respond. However, these seem to be the sorts of properties that defenders of the radical claim have in mind.

aspects while preserving many of their benefits. Thus, we shouldn't join the condemners of metaphor, like Hobbes and Locke, in damning aspects as well.

4.7: ***The Analogy with Aspectual Perception Revisited***

At this point in the chapter, I have laid out my main claims about aspectual thought. In this section I review the most important points from our survey of aspectual perception in chapter 2, and show how my account explains analogous points about aspectual thought.

The discussion in this section should suggest some ways in which the account of aspectual thought might in turn be recruited to an account of aspectual perception. This would be a nice benefit. However, perception poses significant additional challenges that won't be directly addressed by an account of aspectual thought. Specifically, perception, at least vision, essentially involves both conceptual *and* spatial structure. Reconciling these two sorts of structure, and explaining the interactions between them, is a significant undertaking of its own.

4.7.1: *Non-Propositionality*

A primary challenge in making sense of the non-propositionality of perceptual experience was explaining the intuition that perception represents the world as being a certain way without reducing that representational content to something quasi-sentential: to the sort of thing that is usually expressed by a single, literally construed sentence. We saw that this involved two main features in the case of perception: first, that concepts play a structuring role in perception, and second, that perception is semantically denser and more replete than many sentences are. My account of aspectual thought incorporates both of these features.

First, I have built a complex structure into characterizations themselves. In the case of thought, possibly unlike perception, the constituents of that structure are often themselves propositions: for instance, the proposition that Juliet is, fittingly, the most beautiful girl in all of Verona. But these constituent propositions are related to one another by different degrees of prominence and depth, and they

183

are linked together by chains of motivation which make some constituents more central than others. These connections between propositions are most easily thought of as structural rather than propositional — although we can devise higher-order propositional representations of them if we so choose. (I return to this point in §6.1.3.) Because characterizations have these structural relations as an essential part, two characterizations containing all the same propositional constituents can still represent their subjects in markedly different ways.

A second sort of alteration in *how* one thinks about something that needn't involve a propositional change, which is in some ways analogous to this structural change and which we also first encountered in our investigation of aspectual perception, is produced by changes in attitudes. As we've seen, part of the work of aspectual thought is typically to transfer some of our attitudes from the governing characterization to the subject characterization. Because attitudes are permeant, they often also induce further structural changes in the subject characterization as well.

The parable of King David and the prophet Nathan clearly demonstrates how both these sorts of differences can be real differences despite being largely non-propositional. Nathan's parable works so well because it first sets David up with an independent characterization that has a starkly differentiated structure of prominence and centrality, and strong affective attitudes. When David makes the analogical connection, the analogy is tight enough that the affective attitudes are transferred over to his own situation. In all this, the primary change in David's way of thinking is not his acquisition of any new belief about how the world in fact is. Rather, the parable primarily reconfigures the relative importance of the facts he already knows, and thereby also makes different emotions appropriate.

Having isolated these two sorts of non-propositional alteration, structural and affective, we are now also in a position to explain the puzzling cases of aspects that seem not to be underwritten by any substantive concept. Recall that in §2.4.1, I mentioned two such sorts of aspects. First, there were unlocalized aspects underwritten by affective attitudes, such as the aspect induced by the soundtrack for a movie scene, or by a mood. We can now say that these aspects result from an affective attitude permeating the whole of one's characterization of an individual, or sometimes of all one's

184

characterizations at a moment, without applying directly to any one constituent feature. These affective attitudes in turn raise the prominence of those features that fit with them, and motivate the fittingness of other, further features. Second, there were what Roger Scruton called 'merely formal' aspects: those which consist entirely of organizing elements relative to one other, as opposed to endowing any of those elements with substantive significance (for instance, with inferential import). His prime example was hearing a sequence of notes as a melody. We can now say that these aspects are those which consist entirely in imposing a certain structure of prominence, centrality and perhaps depth on a characterization.

The second crucial difference that we isolated between perceptual experience and the propositional content of a typical sentence was that perception is semantically dense and replete. That is, perceptual experiences are both more precise (they can represent a feature as falling at any point along a (nearly) continuous scale) and richer (they represent many features simultaneously) than the contents of sentences typically are. Aspectual thought is likewise much more precise and rich than any one propositional thought, although not *as* precise or rich as perception. Part of this precision and richness is directly inherited from characterizations. Because characterizations often contain highly specific, experientially represented features (such as memories of the subject's particular gestures and attributions of particular emotions to them), aspects in turn often match or introduce such precise, highly specific features. And because characterizations typically contain so many features, aspects typically work by matching many features to one another.

In addition, aspects introduce a sort of richness above and beyond that which is already inherent in characterizations. Although we pursue aspect-application just to the point at which the results justify the effort expended, there are usually many more potential matches than we actually make. Further, we saw in §4.3 that the process of feature matching can be recursive: analogical and especially metaphorical matches can require applying subsidiary aspects, which are more local, less complex, and less recherché, within the overall aspect. These factors make aspect-application indefinitely **open-ended** in practice if not in principle.

185

4.7.2: *Twofoldness*

This last point, about the in-practice open-endedness of aspect application, puts us in a position to explain a third major parallel between perception and aspectual thought. When we look at pictures, we see a scene *in* the surface, while simultaneously being aware of the flat surface itself. This simultaneous awareness of surface and scene is what Wollheim calls 'twofoldness'. The analogous effect in the case of aspectual thought, and in turn of metaphor specifically, is our awareness that we are thinking about the subject *through* or *by way of* the governing characterization.

My account of aspectual thought explains this in the following way. Even as the aspect alters the subject characterization, we continue to be aware of the ways in which the governing characterization differs from the subject characterization. The gap between a and F is especially great, and the feeling of twofoldness is especially strong, when the aspect is an external one. More specifically, in applying an aspect we try to shape the subject characterization in the governing characterization's image as fully as possible. But the subject characterization never becomes fully identical to the governing characterization. (Again, this is increasingly true as the two characterizations represent increasingly different sorts of thing.) First, we match as many features as we can; but inevitably many features in each characterization go unmatched. Second, we match features in the governing characterization to those features in the subject characterization that are most similar to them; but as the two sorts of thing characterized become more different, matches of identity will be increasingly rare. Even matches of scalar similarity involve a significant difference between the two features. Third, we raise the matched features in the subject characterization to a level of prominence and centrality that maximally approximates the structural role of the matching features in the governing characterization; but those features of the subject can still only rarely be assigned the very same structural role as the matching features play in the governing characterization. Because of all these differences, even as the subject characterization is molded in the governing characterization's image as fully as possible, the governing characterization continues to tug harder at the subject characterization, reminding us of the remaining differences between them.

186

4.7.3: *The Two Axes*

The fourth and final parallel I want to take up between aspectual perception and thought is that of Wollheim's two axes. Recall from §2.3.3 that perceptual experiences can vary in two distinct ways, according to the role a concept plays in the experience. First, the concept can require more or less *effort* to apply. Second, it can be applied with more or less *commitment* to the truth of its applicability. I also pointed out, with respect to the second axis, that even when there is no commitment at all to the truth of the concept's applicability to the thing seen or thought about, still there must be some commitment to the concept's appropriateness or fittingness for the thing. In turning to the question of how and why these two sorts of variability manifest themselves in aspectual thought, we also turn from the applicability of a concept as an aspect for a perceived object to the applicability of a governing characterization, F, as an aspect for a characterized subject, a.

Variability in an aspect's ease or immediacy of application is easier to explain. At one extreme, we have aspects that have governed our thinking about an individual or kind for as long as we can recall, but without our ever having explicitly reflected upon them. This is analogous to the way in which we normally always already see a knife and fork *as* a knife and fork. These sorts of aspects are rare, simply because we usually acquire a disorganized jumble of thoughts about an individual over time, rather than acquiring the sort of overall structure that an aspect imposes. Next out along the axis come aspects which strike us immediately upon our first encounter with the subject, and organize our characterization of it from that first impression. This is followed, further out, by aspects that take hold as soon as they are suggested to us, and from there on out by aspects that require increasing amounts of active and conscious effort to apply.

Now, what is the source of this variability? As we saw in §4.4, aspects work successfully to the extent that tight a-matches can be found to prominent F-features without too much effort. In general, this requires that the matching a-features themselves already be relatively prominent. In particular, such matches will be easier to uncover to the extent that there are just a few highly prominent a-features, and these features get matched. Too many potential matches can cloud the aspect's application by dividing

187

our attention among them. I think this is why the most immediately forceful — and persistent — aspects are often those which pick up on just one or a couple of prominent features and match them tightly, such as when I think of someone as a twitching nervous rat because of his facial tics and rapid sidelong glances. This may also explain why it is often easier to apply an aspect for thinking about something with which we aren't too familiar. Because our characterizations are richer for individuals and kinds that are more familiar to us, there are less likely to be just a few highly prominent features in our characterizations of them.

When the aspect's application does require more active effort, this will be because the subject characterization's constituent elements and structure are not already in a form amenable to being governed by F. This can happen for one of three reasons: none of the most prominent elements in our F-characterization may find tight matches to prominent a-features; many of the prominent and central elements in our a-characterization may be actively out of character for something that's F; or the structural relations between those elements that do find matches may be quite different in each characterization.

The second axis of variability concerns our commitment to the aspect's applicability; this itself encompasses several different variables. First, one might think that a really is (or is not) F; one might also be unsure but more or less inclined to think it is. This is the variability that Wollheim initially discussed in the context of seeing-as: whether one thinks the concept truly applies to the object. We can now add that one also might or might not think that F is *fitting* for a, whether or not one also thinks that it is true of it. If one doesn't think it is fitting, then the aspect will be external, in the sense discussed in §4.5.1. Employing *the sun* as an aspect for thinking about Juliet is a paradigmatic case of an external aspect.

However, notice that one can still be fully committed to the applicability of an aspect even if it is external. Indeed, I think Romeo is fully committed to the applicability of the aspect of the sun for thinking about his beloved. When one is fully committed to an aspect's applicability in this way, one thinks that it accurately reflects the subject's true character, by imposing an appropriate overall structure

188

on our characterization of that subject, even though the governing feature itself — in this case, being the sun – obviously does not itself apply.

One can also take a weaker sort of commitment toward the aspect, which I mentioned briefly in §4.6.1: one might be *instrumentally* committed to its usefulness for thinking about the subject. That is, one might think that even though F is neither true of nor fitting for a, and even though thinking of a under the aspect of F produces a distorted overall characterization of a, nonetheless certain important features of a are most clearly revealed by thinking of it under the F-aspect. This is the most common way in which we employ aspects: temporarily, in order to achieve a localized effect. For instance, one might 'try on' a new aspect in theoretical contemplation to see what it turns up, as when scientists first explored the orbital model of the electron. Or, one might decide in a conversational context that the aspect provides the most efficient way to bring attention to or to introduce a certain feature into the characterization of one's interlocutor. For instance, one might describe George as 'a real Einstein' simply in order to bring out the particular sort of ingenious brilliance he displays, but not to apply any of Einstein's further prominent and central features, such as his eccentricity and disheveled appearance. This distinction between the sorts of commitment one can have to aspects themselves will become important in the context of metaphorical content and of paraphrase, in §5.3.4 and §6.1.

4.7.4: *Out on One Axis: Two Different Imaginative Projects*

At the far end of Wollheim's 'axis of commitment' for aspectual perception came 'mere make-believe'. As we've just seen, a similar point goes for thought: we can 'try on' an aspect in thought just to see what its effects might be, without any antecedent commitment to its applicability, even without yet believing that anything new about a's character will necessarily be revealed as a result. We might try the aspect on for sheer amusement, as the mental analogue of seeing scenes in passing clouds. Or we might try it on for some more serious cognitive purpose, such as to understand someone else's point of view on the subject at hand.

For some such far-out aspects, though, even just 'trying them on' requires considerable imaginative effort. These are the aspects I want to discuss in this section. I think they are worth discussing because they represent a pattern of thought in which we do sometimes engage, including in the service of metaphorical comprehension. However, I suspect that these aspects go beyond what is required to explain paradigm cases of either successful aspect application or metaphorical comprehension. For the sort of aspect-application I want to discuss here, the aspect initially does not in fact 'work': we cannot initially execute *either* the first or the second stage of aspect-application — first because the governing concept is too alien to be brought within the subject characterization, and second because there aren't just enough *a*-features available that can be appropriately matched to *F*-features. Thus, these cases are even more extreme than external aspects like *the sun* as applied to Juliet. Despite all these obstacles, we may still eventually succeed in applying the aspect, by way of *first* using our imagination to alter our characterization of *a* so that the aspect *F* can *then* be applied to it.

To get clearer about what this sort of imaginative alteration involves, we need to distinguish two rather different such projects in which we might engage.[41] Distinguishing these projects will also, I think, provide the beginnings of a distinction between the patterns of imaginative thought that are central for understanding metaphor and those that are central for understanding fiction — although I will not be able to pursue that distinction in any depth here.

In the first imaginative project, we imagine that the subject undergoes a transformation which results in its having the character fixed by the aspect in question. Suppose I am trying to think of my highly ebullient and energetic friend Jane under the aspect of retiring placidity. This feature is sufficiently alien to my characterization of her that I can't use it as an aspect for thinking about her. I might eventually succeed in applying the aspect, though, by imagining her in the future going through a painful divorce or a death in the family, which then strips away her ebullience and replaces it with a sadder, perhaps wiser calmness. In order for me to be thinking of *her* through this process — of Jane as

---

[41] Discussions with Richard Wollheim were crucial in getting clear on this distinction.

190

opposed to some mere counterpart — it's important that my understanding of her actual current character play a central role in guiding her imagined transformation. But it's also possible, through this sort of imaginative process, for me to imagine her character changing quite dramatically.

In this case, I transform Jane in my imagination so that imagined-future-Jane really is retiring and placid; I then organize my characterization of imagined-future-Jane as a whole under this aspect. This case is one of applying a literal, internal aspect. I might, though, also use this same sort of process to apply an external aspect. Thus, suppose I am asked to think of Juliet as the sun, but I am so convinced that she's no more than a pallid, pasty little brat that I can't apply the aspect to her at all. I might then imagine her being sent away to school in Venice, picking up some big city sophistication, and gaining some experience with the darker side of life which makes her more caring toward others. I might thereby be able to think of imagined-future-Juliet, now returned to small-town Verona, as the sun. Of course, without a current characterization to go on in seeking matches, it will be much more difficult for me to figure out how to transform the subject so that a still-merely-external aspect would apply. But it *is* sometimes possible, especially given appropriate contextual guidance about what sort of individual the aspect is supposed to apply to. In a conversational context, such guidance is likely to be present.

In the second imaginative project, we imagine how the subject's *current* character, rather than its imagined-future-character, would need to be different if *F* were to be applicable as an aspect to that character. We accomplish this by imagining some larger context of *other* features, which are not now part of our characterization of *a*, but which, *if* true of and fitting for *a,* would make *F* itself applicable. Thus, for instance, I might succeed in thinking of Jane under the aspect of retiring placidity by reconceiving her ebullience as a protective defense, built up to enable her to engage with a world she finds terrifying. In this case, it might help me in applying the aspect to imagine her collapsing at home at the end of the day, turning off the phone's ringer, happy to be alone for a few hours before social obligation calls once more.

Again, a similar point also goes for an external aspect: asked to think of Juliet as the sun, I am balked, for the same reasons. But I might eventually succeed by way of imagining that what I take to be

191

Juliet's brattiness is really a sort of self-determination; that she seems pallid only because she is focused on distant, lofty matters like philosophy; that her apparent pastiness is really a pure white-hot glow of intense energy. *Then* I may succeed in applying the aspect to her. When I reconceive a subject's characterization in this way, I imagine that some of my beliefs about its character are faulty, and I try imaginatively to save as much as possible of the character's appearances while reconceiving the underlying reality as needed to apply the aspect.

These two ways of using the imagination to reconceive a characterization do seem to be rather different. In particular, we can imagine something's being transformed in the future so as to become capable of having a property (in the first project) that we cannot imagine it as currently having (in the second). So, for instance, I might find no way to make sense of Alice's *now* being a contented stay-at-home mom, but I may be able, with relative ease, to make sense of her becoming such a person. I am primarily interested in the second imaginative project, because I am trying to explain what it means to see or think of some thing under an aspect, where that means achieving insight into how that thing *is*. This is, I think, the typical primary aim of metaphorical communication: to make claims (and other speech acts) about how things *are*, sometimes by stretching our imaginative powers in directions we hadn't expected them to go. By contrast, I think that the main aim of fictional discourse is to help us imagine other ways in which things *might* be; as such, it seems to be primarily concerned with the first imaginative project.

It would obviously take significantly more work to articulate this claim about the distinctive imaginative projects of fiction and metaphor fully, let alone to make it convincing. However, I think we can see enough here to rebut an objection that is sometimes made against certain theories of metaphorical interpretation. Eva Kittay, for instance, criticizes Samuel Levin and other lexical semanticists on the grounds that they can't distinguish metaphor from fairy tales.[42] In metaphor, according to these semanticists, we suspend certain restrictions on the application of certain terms, such as 'is a pig'. But,

---

[42] *Metaphor*, pp. 159-161.

the objection goes, this seems to be just what happens in fanciful fictions as well. So, for instance, in interpreting both

    (5) Sam is a pig,

intended metaphorically, and

    (6) Odysseus was a pig,

intended as a claim about an episode in the *Odyssey*, it seems that we suspend the assumptions that humans can't be pigs, that pigs can't talk or reason, and so on. What's the difference between the two interpretive processes supposed to be? With just the usual lexical semantic machinery of deleting and transferring selection restrictions, there is no more to say.

       I think I can offer such a difference — even though my account of metaphorical comprehension *per se* is not yet on the table. In interpreting a sentence in fiction, we imagine that the subject really is such that the predicate in question applies to it. This often requires transforming what we pretend to be the subject's character traits in many other ways as well: so, for instance, when I imagine that Tatanya, the näive, dreamy country girl in *Eugene Onegin*, has married an older, urbane Prince, I also imagine that her patterns of speech, carriage, and dress, and perhaps her desires and goals, have changed as well. These imagined transformations pile on as the narrative passes — indeed, the narrative can be treated as a series of instructions for imagined transformations (though I'm not sure this is the best analysis of fiction). Now, none of this typically involves applying the relevant predicate's associated characterization as an aspect. Sometimes it does involve aspect-application, as when a particular described action of the subject is supposed to be central to or emblematic of her entire personality (imagine a fiction in which a character, Carol, joins a Satanic cult; or one in which our hero, a patriotic New Deal Democrat, becomes a Soviet spy...). In such a case, the aspect applies in just the way it would in the sorts of purportedly true situations with which I opened this chapter: we are supposed to add the relevant feature itself to our characterization in a central and prominent way, and reorganize our characterization as much as possible in its image.

193

In interpreting metaphor, by contrast, we are much more likely to encounter an external aspect, one for which we cannot bring the relevant feature within our characterization of the subject. Even when the aspect is internal, we are not supposed to pretend that the relevant feature really is true of the subject, as we do with fiction, but rather to understand how it would be fitting *if* it were, and to determine the speaker's intended content in the manner I will describe in chapter 5. Thus, the overall cognitive effect of the aspect is different even when the same sort of aspect is involved. Finally, recall also that *most* metaphors don't require us to engage in the second sort of imaginative project at all, for the reasons cited above: there usually are enough matches between the two characterizations as they stand for the aspect to work, and we are usually being asked to cultivate the aspect for thinking about subject as it now *is*. Thus, the first imaginative project, of imagining ongoing transformations, is much more central to our engagement with fiction than the second is to our engagement with metaphors.

4.8: *Conclusion*

The cases I've been discussing in this last section lie, by definition, at the extreme outer limit of aspectual thought. Thinking about them shouldn't draw too much attention away from the much more common, and much more comprehensible, sorts of examples with which I opened this chapter. In this chapter, I have explained how applying an aspect brings about changes in a subject characterization's structure, and sometimes in its constituents as well. These changes produce real differences in how the subject is represented; these differences are real although often most naturally thought of as non-propositional.

Aspects accomplish this effect by molding the subject characterization as fully as possible in the image of the governing characterization. In applying an aspect, we seek matches in the subject characterization to the most prominent and central features in the governing characterization, and then raise the prominence and centrality of the matching *a*-feature to maximally approximate those of the matched *F*-feature. Matching can be made in virtue of the identity of features, or of their scalar,

analogical, or metaphorical similarity. Under appropriate circumstances, features can also be introduced into the subject characterization.

In chapter 5, I exploit the account of aspectual thought developed in the last three chapters to explain metaphorical comprehension and content. In chapter 6, I take up the vexed question of whether metaphors can be paraphrased, and I bring out some implications of my investigation into metaphor and aspects for thinking about thought more generally.

CHAPTER 5: THE FORCE AND CONTENT OF METAPHORICAL UTTERANCES

In this chapter I return to metaphor *per se*, and show how my theory accounts for the content and comprehension of metaphorical utterances. Metaphors accomplish two main types of effect. On the one hand, we saw in §1.2 that metaphorical utterances are indeed speech acts: by speaking metaphorically, speakers commit themselves to certain conditions obtaining, undertake obligations to or direct their hearers to make those conditions obtain, and so on. Existing semantic and pragmatic accounts recognize this fact, but lack the resources to explain adequately how the relevant conditions could be determined on the basis of the literal meanings and modes of combination of the words uttered. On the other hand, we saw in §1.3 that metaphorical utterances also essentially involve a sort of non-propositional understanding, of making us 'see things in a new light', or under a new aspect.

Both pragmatists and non-cognitivists have tended to assume from the outset that no single theory can explain the full range of metaphor. They both assume, that is, only some metaphorical utterances serve to undertake speech acts with genuine truth-conditions, that only some serve to induce non-propositional effects, and that the two classes of metaphor are largely exclusive. I rejected this assumption, and proposed to develop a more encompassing, unified theory of metaphor. Specifically, I proposed that metaphorical utterances exploit non-propositional aspectual thought as the means to a cognitive, propositional end, rather than as the sole end in itself. In order to for this proposal to be a substantive one, I needed to develop a theory of aspectual thought itself. Armed with the theory developed in the last three chapters, I can now show how aspectual thought serves the end of communicating propositional content.

After laying out my basic account in §5.1, I extend my account in §5.2 to cover metaphorical utterances in which the aspect cannot simply be read off from the sentence uttered. In §5.3, I discuss some of the most important variables among metaphors, specifically those concerned with the relationship between the aspect cultivated and the content intended, and show how these variables conspire to produce the difference between 'ordinary' and 'poetic' metaphors. Finally, in §5.4, I contrast my account with

those of several other philosophers which we have discussed along the way: of Black, Searle, Davidson, Hills, and Stern, among others.

5.1: *The Basic Story*

Here, in a nutshell, is my theory of metaphorical comprehension. A metaphorical utterance is an utterance in which the speaker does not primarily intend to commit herself to what she says. Instead, she intends for her hearer to take her to be committed to a propositional content which is determined by means of cultivating an aspect in thought. Metaphor is thus distinguished from literal utterances that also generate aspects, such as the literal, true or apocryphal stories with which I began chapter 4, in that for a metaphorical utterance but not a literal one, the aspect serves as the means for identifying the speaker's primary intended content. Metaphor is distinguished from other indirect speech acts in that in metaphor, aspectual thought serves as the function taking us from what is said to what is meant.

I had already claimed this much at the end of chapter 1. Now, though, I can be much more specific about just how the speaker's metaphorical meaning is determined. In the simplest case, the speaker says '*a* is *F*', and intends her hearer to think of *a* under the aspect of *F*. By applying *F* as an aspect for thinking about *a*, the hearer can identify those features of *a*'s now-reconfigured characterization which are most tightly matched to the most prominent *F*-features. These, I claim, are the features that the speaker predicates of *a*. As a bit of shorthand, I'll call them *M*-features.[1]

Let me take as an example the familiar case of Romeo's utterance of

(1) Juliet is the sun.

In order to determine Romeo's intended *M*-features, we determine which features in our characterization of Juliet are most tightly matched to the most prominent features in the sun's characterization. I've already mapped out the most important matches in Figure 6, in §4.5.4 (repeated here). The *M*-features, then, are that Juliet is the most beautiful lady in Verona, that she is an exemplar of goodness, that she

---

[1] When the matched constituents are affective attitudes rather than properties, the speaker predicates *it is fitting to feel* A *in relation to* a.

FIGURE 6 REPEATED HERE

nurtures Romeo's maturity, that she makes Romeo feel emotional and physical warmth, and that Romeo's thoughts begin and end with her. These features don't by any means exhaust the *effects* Romeo hopes to induce in his hearers. Nor do they even exhaust what Romeo wants to suggest is *true* of Juliet. But — *modulo* contextual factors to be discussed below — they do form the core of what he commits himself to by his utterance.

This fixes the metaphorical utterance's intended *content*. Its illocutionary force is in the first instance broadly determined by the grammatical mood of the sentence uttered. Once the metaphorical content has been determined, the hearer uses this, along with other conversational clues, as a guide for identifying the specific intended illocutionary act.[2] Thus, for instance, the syntax of

(2) Go out there floating like a butterfly and stinging like a bee

determines that it must be a directive; the fact that it is an order rather than a warning or a recommendation is indicated both by the specific metaphorical content and by the general conversational context — for instance, by the power dynamics between the speaker (a coach) and the hearer (an athlete).[3]

Returning now to content: to understand how aspectual thought can determine the appropriate content in every case, we must attend to the fact that in the context of a metaphorical utterance aspectual thought is playing a primarily social, communicative role, rather than a primarily individual, contemplative one. As I first discussed in §3.7.2, characterizations are quite contextually malleable. In particular, within a given conversational context the hearer will, insofar as possible, adopt what he takes to be the *speaker*'s characterizations of the relevant individuals and kinds. (In addition, the speaker may provisionally alter her own characterizations to accommodate her hearers'.) Further, the specific way in which the hearer applies the aspect — which matches he makes and in which order — will be influenced

---

[2] See Bach and Harnish, pp. 34-37.
[3] Notice here that metaphor, like figurative language generally, is distinguished from implicature in that the force of the speaker's intended speech act is directly inherited from the grammatical mood of the sentence uttered. With non-figurative indirect speech acts, one can mix forces. For instance, one can use an assertive sentence to make an indirect request: an utterance of
(1) You're stepping on my toe
might serve to request that the hearer change his stance. (Cf. Searle, "Indirect Speech Acts.")

by his understanding of the conversational context, and specifically of the speaker's communicative intentions.

This **adaptive process** applies to communication across the board, of course, but it is more complex and pervasive in the metaphorical case. When the differences between the speaker's and hearer's characterizations are relatively minimal, the hearer's adaptation may be relatively unconscious, and consist mostly of realigning the prominence of particular features within his characterizations. More significant differences require more explicit interpretive effort, and may eventually undermine the possibility of communication altogether. Nonetheless, the cooperative nature of conversation is so fundamental that, so long as the hearer can figure out what the speaker's characterizations are, then he will adopt them for the purposes of comprehension — even if he strongly disagrees with them, and even if the two interlocutors are at conversational odds, as in a courtroom interrogation.

The key role that conversational context plays in determining how the hearer cultivates the aspect in turn explains at least three further features of metaphorical comprehension and metaphorical content. These are, briefly, differences in the aspects generated by different utterances of the same sentence, ease of feature introduction, and variations in content produced by variation in conversational weight.

The first feature that is explained by the communicative context is that in interpreting different speakers, the hearer may cultivate different aspects, and so identify different contents as intended. At a minimum, the hearer must grasp the speaker's 'essential description' of the subject being described, in the sense specified in §4.5.3. Different speakers may have very different essential descriptions of the same subject. So, for instance, the aspect under which Juliet is cast when Romeo utters (1) is very different from the aspect that would be generated if Benvolio had uttered the same sentence in the course of warning Romeo away from the dangers of loving her. The conversational context, and even just the intonation, of their respective utterances make clear that they characterize Juliet in very different ways. For Benvolio, Juliet's most prominent and central feature by far is being a Capulet, and thereby being a deadly enemy of the Montagues. Thus, if Benvolio were to utter (1), the hearer would appropriately take the sun's distance to be matched by Juliet's being socially untouchable, and its potential to burn to be

matched by her potential to destroy Romeo's life. As we saw in §1.1.2.2, Benvolio can exploit these potential matches, which Romeo must ignore, in order to produce a very different metaphorical meaning.

In a normal conversational context, when contextual adaptation fails to bring the hearer's characterizations fully in line with the speaker's, the only matches that a hearer can appropriately exploit in determining the speaker's meaning are matches between features in the speaker's characterizations. For instance, because the sun's being the center of the solar system is *such* a prominent and central feature of my sun-characterization, it is nearly irresistible for me to match the sun's location at the center of the solar system to Juliet's being the center of Romeo's world. It just fits so well for me with the rest of Romeo's characterization of his beloved. But because that feature of the sun was not an element of Romeo's characterization of the sun, that feature of Juliet cannot be part of Romeo's metaphorical meaning. However, a speaker *can* end up committing herself to a feature she didn't explicitly intend, so long as that feature is included as fitting within her subject characterization and can be appropriately matched to a feature in her governing characterization.

The second important effect of the conversational context on the aspect is that it greatly facilitates feature introduction. The hearer assumes that the speaker's utterance is motivated by the desire to communicate some more or less determinate content. He combines this assumption, and more specific presumptions about what sort of speech act the speaker might be making at this point, with the aspect. If the aspect cultivated doesn't itself already provide *M*-features that the speaker could plausibly be predicating of the subject in this context, then the hearer will seek to introduce such features into the subject's characterization in the manner described in §4.3. As a result, metaphors can indeed be, and often are, genuinely informative, and don't just draw our attention to already known but neglected features.

The way in which conversational context and aspect-cultivation interact to facilitate feature introduction explains how we can comprehend even metaphors for which we don't have particularly rich subject characterizations. We saw in §1.2 that Beardsley objects to a 'comparison theory' of metaphor on the grounds that it makes truly informative metaphors impossible:

> Suppose the poet remarks, 'My sweetheart is my Schopenhauer.' According to the comparison theory we are to ask what his sweetheart and Schopenhauer have in common. But we don't *know* his sweetheart, so how can we answer this question until he tells us, by the metaphor itself, what she is like?[4]

I can now respond to this objection. Given an appropriate governing characterization, and armed with an essential description of the subject, we generally do have the resources needed to generate the most important matches between the two characterizations, and thereby to isolate the primary contents of the metaphorical speech act. Thus, suppose I don't know anything about Romeo's characterization of Juliet. Nonetheless, I can still discern from his tone and the context that he intends the utterance as a metaphorical compliment, that he is a lovelorn young man, and that Juliet is the object of his affection. From this, I can generate the essential description 'girl who's his beloved', and with it at least a minimal characterization of Juliet. This puts me in a position to isolate at least the most important contents of Romeo's claim, by introducing the relevant features as necessary. The more I know about his characterization, the richer my understanding can be, but even a relatively minimal grasp of his basic way of thinking about Juliet will enable me to interpret his utterance. Absent this grasp, however, I will not be able to understand him.

I confess that I *don't* actually understand Beardsley's poet's metaphor, but the difficulty here is not that I don't know his sweetheart, but that I don't have any sense for the appropriate governing characterization of *Schopenhauer*. By contrast, I do have a sufficiently clear sense of the common sun-characterization to get going on Romeo's utterance. But, as this contrast between Romeo's and Beardsley's poet's utterances reveals, metaphors that rely on more novel aspects require more contextual support to underwrite feature introduction. If the aspect is quite novel, we will have little sense for what sorts of as-yet unknown features the speaker might intend for us add to our subject characterization as matches to features of the governing characterization.

The third important consequence of the conversational context is that the effort the hearer puts into cultivating the aspect will depend on the conversational weight that the speaker accords to her utterance. If the speaker offers the metaphor as a mere off-hand remark, the hearer will only seek to

---

[4] "The Metaphorical Twist," p. 266.

identify the one or two most prominent *M*-features. When a metaphor is given greater conversational weight, the speaker expects for her hearer to take the aspect more seriously, and thereby to identify a fuller range of features. The greater the metaphor's conversational weight, then, the more open-ended and indeterminate its content will generally be. Metaphors in poetic contexts, such as Romeo's utterance, are prime examples of this. Thus, although the *M*-features I specified for (1) above do form the *core* of Romeo's intended claim, hearers are permitted to include additional appropriately matched *a*-features as well. I return to conversational weight in §5.3.1.

5.2*: From Utterance to Aspect: Complexities*

So far, I have specified how the content (and force) of a metaphorical utterance is determined for the basic case. This is the case in which the speaker says '*a* is *F*', intends for her hearer to cultivate a characterization of the kind *F* as an aspect for thinking about the individual *a*, and thereby isolates a set of *M*-features which she intends to predicate of *a*. In such a case, 'is *F*' is construed metaphorically, while '*a*' designates its subject by literal means, in the simplest instance by naming it. Not all cases fit this model, however. In order to apply my account more generally, I now need to address two sources of complexity. First, not all sentences used metaphorically take this simple '*a* is *F*' form. Terms in nearly all parts of speech, complex phrases, and even whole sentences can all be used metaphorically. Second, the aspect generated cannot always be 'read off' so directly from the sentence uttered: it may only be indirectly implied by it.

5.2.1: *Syntactic Complexity*

As I said in §4.5.2, there is a loose sense in which my account is at root predicational: an aspect always involves two characterizations, and the governing characterization is applied *to* the subject characterization. It is because of this quasi-predicational structure that the utterances of the basic '*a* is *F*' form are basic: they most clearly reveal how metaphor works. However, governing as I've described it is obviously quite different from predicating. In particular, no one feature need be introduced into the

subject by the aspect's application; the alteration in a characterization may be entirely structural. More importantly, as I emphasized in §4.5.1, characterizations themselves have basically the same type of constituents and structure whether they are of individuals or of kinds. As a result, aspects may employ any combination of characterizations of individuals and kinds in both governing and subject positions. More generally, as I said in §3.7.1, although characterizations are often closely, even conventionally associated with words, there is no essential connection between the two. We have characterizations of individuals to whom we can refer only demonstratively or with a name, and of kinds we can only describe with a complex phrase.

For all these reasons, different syntactic structures for the sentences uttered don't generate fundamentally different kinds of aspect. Because on my account, the process of cultivating the aspect occurs as part of the *pragmatic background* to comprehending the metaphorical meaning of the sentence uttered, my account has the flexibility required to deal with metaphorical utterances of sentences with very different syntactic structures in a unified way. However, even if there is no essential difference among the aspects generated by syntactically distinct sentences, I still need to make clear how we know which aspect to generate on the basis of the sentence uttered. Further, the logical form of the speaker's meaning *is* importantly constrained by the syntactic structure of the sentence uttered, as it is not so constrained for implicature. I need to explain how this constraint works as well.

As I said, in the simplest case an individual, *a,* is referred to by a literally applicable name or demonstrative, as with

(1)   Juliet is the sun.
(3) That guy is a pig.

Cases where the subject is an individual referred to by a literally-applicable description work in just the same way:

(4) My beloved is the sun.
(5) The undergraduate who left this room in this sorry state is a pig and should be housed in a barn.
(6) That gentleman in the tattered suit with the quiet voice is in fact a bulldozer.

In these cases, the subject characterization is still of an individual even if the syntactic structure of the sentence is quantificational. The descriptive material supplied in the noun phrase may itself inform the intended subject characterization, as in (4) and (5). In (5), the noun phrase generates a characterization of an individual, but it is of that individual, whoever he or she may be, simply in virtue of him or her fulfilling the noun phrase's descriptive content. Alternatively, the noun phrase's descriptive material may be irrelevant to or even at odds with the intended subject characterization, as in (6). (Here I'm imagining that the contextually salient gentleman *appears* to be a pushover, and that the noun phrase reflects this.)

In §4.5, I discussed aspects in which a characterization of a kind governs a characterization of a kind, such as

(7) Surgeons are butchers.
(8) Some butchers are surgeons.
(9) No man is an island.
(10) Most men are rapacious wolves.
(11) Man is a necrotic, invasive cancer on the earth.
(12) The work is the death mask of its conception.[5]

In these cases, the aspect can still be 'read off' directly from the sentence uttered, and the *M*-features are still determined and predicated of the subject in just the same way as the basic cases discussed in §5.1. ==A case like (8) or (10) introduces one additional interpretive step, because of the **quantifier**.== The aspect itself in (8) simply applies our characterization of surgeons to our characterization of butchers in general, much as (7) does in reverse; the quantificational force is added back to the metaphorical interpretation of (8) after the *M*-features are determined. I believe (9) also works in the same way: first we apply our island-characterization to our characterization of humans, and then we take the speaker to be claiming that no man has the features thereby identified.[6]

---

[5] Walter Benjamin, "The Writer's Technique in Thirteen Theses."

[6] Thus, I am committed to the claim that the *M*-features of (9) are the same as the *M*-features for a typical utterance of
    (9') Every man is an island.
I believe this to be true; indeed, it is a prime piece of evidence for my claim that ==the negation in (9) is external to the metaphorical meaning.== However, it also might be the case that the different contexts in (9) and (9') make different subsidiary *M*-features more plausible. For instance, it might be more easily compatible with (9') than with (9) that there is some travel between islands, and so some interaction between men. I think that in both the negated and the universally quantified cases, this match was already there to be made, and so, for instance, that a hearer might respond to either utterance by making that match explicit. But to point this out would weaken the point of an

This same interpretive pattern applies to all cases in which the syntactic structure of the metaphorical sentence is complex, and so cannot mirror the simple quasi-predicational structure of the aspect itself. In every case, the hearer identifies the aspect, uses it to isolate the appropriate *M*-features, and then feeds those features back into the more complex syntactic structure in place of the metaphorical term. That new metaphorical content then combines with the contents of the non-metaphorical constituents of the sentence uttered to determine the metaphorical meaning of the whole utterance.

In both (10) and (11), the verb phrase's own additional descriptive material adds some extra complexity to the aspect itself. The adjectives ('rapacious' in [10] and 'necrotic, invasive' in [11]) make the governing characterizations more precise, and its associated affective attitudes more starkly negative. In (11) and (12), though, the verb phrase's embedded prepositional phrase doesn't work to alter the governing characterization *per se*. Rather, it guides the aspect's application, by mentioning in literal terms a feature in the subject characterization which should be matched analogically to a central but unmentioned feature of the governing characterization. Thus, it's a central and prominent fact about cancers that they grow on bodies; 'on the earth' makes clear that *body* in the governing characterization should be matched to *earth*, which is what man lives on or inhabits. This match then suggests a host of subsidiary matches, for instance between the uncontrolled exploding growth of cancers and of man's population, and between the cancer's consumption of the body's tissues and nutrients and man's consumption of the earth's trees and minerals. In a similar fashion, death masks do capture the look of a

---

utterance of (9) more than it would an utterance of (9'), and so would be less compatible with a sympathetic reading of (9). Thanks to Dmitri Tymoczko for raising this point.

David Hills has cited (in conversation) a passage from Wallace Stevens' "The Idea of Order at Key West" to make the same initial point, that negation is external to the metaphorical meaning:

The sea was not a mask. No more was she.
The song and water were not medleyed sound
Even if what she sang was what she heard,
Since what she sang was uttered word by word
It may be that in all her phrases stirred
The grinding water and the gasping wind;
But it was she and not the sea we heard.

In this passage, the poet is rejecting the claim that the sea is a mask through which a fundamentally anthropomorphic emotion expresses itself, and likewise rejecting the claim that 'she' is a mask through which the sea expresses its distinctive nonhuman cry, even as he acknowledges some of the more basic facts which might motivate this metaphorical characterization.

memorialized person's face at the moment of death, but are themselves a mere shadow of the person's vivacity while alive. Works of art by analogy do capture the state of their creators' mind at the moment of conception, but reveal only a small portion of their creators' creativity. (Notice that while we could add 'some' to (12) to make a claim with a different quantificational force, as with (8), (9), or (10), this is not possible for (11), because the metaphor's central analogy matches *man* to a single undifferentiated mass. Thus, the aspect itself can set important constraints on which sentences can be used metaphorically.)[7]

5.2.2: *Metaphorical Noun Phrases*

So far, all of these examples have simply added a few complications to the basic pattern. The subject is still referred to in literal terms by the noun phrase, and the verb phrase still directly introduces the aspect's governing characterization. In the sort of case I now want to discuss, the subject itself is specified metaphorically through the noun phrase, and the verb phrase is construed literally. In order for such metaphors to be interpretable at all, there must be some clear contextual clues for identifying the subject under discussion. In particular, either a certain object or person must already be independently salient, or else the metaphorical description must be sufficiently routinized to give the hearer a grip on what sort of features the speaker must intend to apply. So, for instance, one might utter

(13) The fox has been fomenting gossip and discord again,

knowing that 'the fox' is often used to describe a person as cunning and tricky, and expecting that some cunning and tricky person (say, George) is sufficiently salient for the hearer to identify him as the subject. If we make explicit this additional step of identifying the subject, we might restate the speaker's intended meaning in terms of an appositional sentence:

---

[7] I dwell on the point about quantification at some length in part because my account is likely to appear overly simple given its fundamentally quasi-predicational structure, and given that a subject/predicate structure is insufficient for an analysis of language generally. The fact that something analogous to predication *is* at work in metaphorical comprehension is supported by the fact that when both the subject and governing characterizations are of kinds, still it is only possible to interpret the predicating phrase metaphorically. We cannot, for instance, employ (7) to claim that butchers are metaphorically surgeons; nor can we employ

(10') Man is a wolf

to claim that wolves are metaphorically manlike. Thanks to John MacFarlane for pressing this point.

(13') George, the fox, has been fomenting gossip and discord again.

The process of metaphorical interpretation focuses on the implicit metaphor

(13'') George is a fox,

which is one of our basic cases accounted for in §5.1. If we take the *M*-features for (13'') to be being

cunning and tricky, then the complete metaphorical content is given by

(13''') George, the cunning and tricky guy, has been fomenting gossip and discord again.

 Thus, the only significant difference between (13) and the basic cases discussed earlier is that the subject

characterization is introduced by independent contextual means. The metaphor is then entirely

compressed into the noun phrase.

5.2.3: *Aspect Introduction by Colligation*

The third sort of case requires a more significant extension of my basic account. In these cases

the subject of the metaphor is once again referred to in a literal manner by the sentence's noun phrase, and

the verb phrase is again metaphorical. But this time the intended *M*-features are not features of the

subject *per se*, but of a particular state or an action of the subject's. Take for example an utterance of

(14) Jane was all wound up.

'Jane' is intended literally; 'was all wound up' is metaphorical. Here it's not just Jane in general who is

being characterized, as Romeo's utterance of (1) characterized Juliet, but rather Jane at a certain moment,

in a certain state.[8] Because characterizations represent individuals and types of individuals rather than

states and actions themselves, the relevant governing characterization cannot be given directly by the verb

phrase. Rather, it must be introduced "by being hooked up to [a] conventional colligate" of the verb

uttered, as Andrew Goatly says.[9] That is, the verb phrase is more or less conventionally associated with a

---

[8] Cf. the distinction in Spanish between the two verbs for 'to be': *ser* and *estar*.

[9] *The Language of Metaphors*, p. 86. 'Colligation' refers to the common co-occurrence of words. Goatly is here
claiming that *imagery* is introduced by colligation; I'm claiming that the aspect which generates the metaphor's
content is thus introduced. Eva Kittay (*Metaphor: Its Cognitive Force and Linguistic Structure)* also appeals to
colligation as an important mechanism for generating semantic fields. My treatment of these colligational cases in
terms of indirectly-introduced characterizations is influenced more generally by Roger White's theory of metaphor
as always involving the juxtaposition of two separate complete sentences, one literal, the other metaphorical; see

certain sort of thing or individual: the sort of individual who can be in the state or undertake the action that the verb phrase explicitly specifies. Given the verb, then, we can generate the relevant, but merely implicit, governing characterization. It is of *that* conventionally associated sort of individual undertaking the action (or being in the state) that is directly and literally specified by the verb phrase. In this case, Jane is to be thought of under the aspect of something that's in the state of being wound up, and one must infer by colligation from the verb that this something is a machine, most likely a clock.

Notice that the inference here about colligation depends at least as much on **mere factual knowledge** about what sorts of things get wound up as on specifically semantic or grammatical knowledge. This is another important reason not to treat metaphorical meaning as a distinctively semantic matter, even though colligation, which plays such a crucial role here, is a phenomenon of words. Even the knowledge of word use which is relevant to metaphorical interpretation may still not be distinctively semantic knowledge, but rather **sociological** knowledge. Those who invoke colligation typically ignore this distinction.

If we now turn to the subject characterization, we find much the same complication. The noun phrase tells us in literal terms who or what the subject is, but the speaker's intended *M*-features are supposed to describe a particular action or state of that subject, rather than that subject itself. We are not just supposed to think of Jane as a clock *simpliciter*, but of Jane as a clock in a particular state: of being wound up. More specifically, we are not supposed to think of Jane *in general* as a wound-up-clock, but of Jane *as she was at that moment.* And just how Jane was at that moment is not explicitly specified by the noun phrase. Thus, in general in cases of this sort, the subject characterization needs to characterize, not just the noun phrase's denotation (as was the case for Romeo's utterance of [1]), but that denotation *being in some as-yet-unspecified state* or undertaking some as-yet-unspecified action. The primary task of the aspect in such a case is to identify just what that action or state is, through feature introduction —

The Structure of Metaphor: The Way the Language of Metaphor Works. White inveighs against predicational accounts like Black's on the grounds that they cannot accommodate metaphors of more complex syntactic forms; this drives him to the complete-sentence analysis. I think my account shows that making a sufficiently strong distinction between metaphorical language and the patterns of thought involved in interpreting it enables an account which is 'predicational' in the way mine is to explain these more complex cases as well.

although as usual, the aspect also thereby casts the subject itself in a more general light. Putting this all together, in the case of (14), the governing characterization is of a wound-up-machine, and the subject characterization is of Jane-in-a-certain-state. Some of the *M*-features are *being under a lot of tension* and *being ready to spring into movement*.

In general, when a component of the governing characterization is introduced by colligation, the governing characterization as a whole, and the resulting *M*-features, will be richer and more specific to the extent that the sort of thing or individual introduced by colligation is likewise more specific. Because characterizations always characterize individuals and kinds of individuals, though, *some* sort of object always must be determined by colligation with the verb phrase. In (14), there isn't a very strong conventional association between 'wound up' and just one sort of machine, although we have the most experience with clocks being wound up. As a result, the metaphor is not that rich or specific. By contrast, in

(15) The earth pirouettes around the sun.

'pirouettes' is conventionally associated specifically with ballet dancers; the action described and the aspect under which the earth is thereby cast are concomitantly more specific. Likewise, in

(16) The gills were kneading quietly.[10]

'kneading' introduces by colligation a quite specific situation for the governing characterization: *hands* knead *dough*. Here, we also have a further complication. Even though the metaphorical sentence itself does not specify a direct object, the verb requires one. And we know as a matter of fact, if not as a matter of semantics, that dough is the sort of thing which typically gets kneaded.[11] By analogy, then, the subject characterization must also include an object of the action. It must characterize, not just gills acting in some quiet way or other, but acting *on something*. We then make the inference that they are most likely acting on water. The specific kind of action the gills are doing is introduced by direct match to the kind

---

[10] Ted Hughes, "Pike," cited by Goatly, p. 86.

[11] The American Heritage Dictionary says that 'knead' is a transitive verb. Perhaps this is wrong, and 'knead' has a separate, genuinely intransitive use as well. If so, then we still know, as a matter of factual knowledge, that when you knead you must knead *something*, which is typically dough. Again, characterizations, and as a result metaphorical meaning and comprehension, are indifferent to this semantic distinction.

of action that hands do to dough when they knead it. Subsidiary matches then include the ongoing repetition of the two actions, and the resistance offered by the dough and the water respectively to the kneading hands and gills.   (The gills might also be acting on air; if so, this would raise to prominence the fact that air is to fish as dough, and water, are to us: a heavy substance requiring effort to manipulate.)

(16) exemplifies an especially common form of aspects introduced by colligation: the verb used metaphorically specifies a *relation* between an actor and an object. Not all such aspects take this form, as (14) and (15) demonstrate; but it is quite common. In (16) the object being acted upon was not mentioned explicitly, and had to be inferred through the application of the aspect. But in

(17) The chair plowed through the discussion.

both the actor and the object are explicitly specified in literal terms. The subject characterization is therefore *the chair doing something to the discussion*, and once again the primary task of the aspect is to determine just what he was doing. The governing characterization is generated by colligation from 'plowed through'. Although this suggests something about the sorts of object doing the plowing and being plowed, 'plowed' is like 'wound up' in being commonly employed — in what appear to be literal terms — in relation to a relatively broad range of agents and objects: ploughs plow through the earth, but ships also plow through the sea, and bulldozers plow through the snow.[12] Because there is such a multiplicity of candidate agents and objects, the governing characterization can legitimately contain only those elements which are common to all of them. What is common to all these plowing agents and their objects is moving in a direct, forceful way that cuts through obstructions without undue attention to what it is moving through. This, then, is what the chair is doing to the discussion. The application of the aspect in turn brings out features of the chair such as his or her unwavering force (and perhaps unreflectiveness or inflexibility), and features of the discussion such as its malleability (and thus in turn, the passivity of the discussion's other participants).

---

[12] Perhaps only the first application is fully literal, and the second two still at least marginally metaphorical. I discuss 'dying' and dead metaphors in §5.3.5.

More generally, to the extent that a verb can be used, like 'wound up' or 'plow', to apply literally in connection with a wide range of types of object, we cannot form as rich, detailed and nuanced a characterization of the particular sort of individual who instantiates the property specified by it, or of the particular sort of situation in which that property is instantiated. As a consequence, there will be fewer potential matches to be cultivated, and the metaphor will be less rich. Further, to the extent that a verb can be used to apply literally in connection with a wide range of types of object, we are less likely to have a strong intuition that that verb is being applied to a sort of object which is inappropriate, given its conventional, literal meaning. As a consequence, the intuition of metaphoricity will tend to be concomitantly diminished. So, for instance, 'pirouettes' applies only to the action of dancers, and so any other application is obviously metaphorical. With a verb like 'cut', which can apply literally in connection to an even broader range of agents and objects — scissors and cloth, mowers and grass, knives and cakes[13] — it is difficult to generate an intuition of metaphoricity at all. 'Plow' falls somewhere in the middle.

When a verb of this sort occurs in an extended context, however, then that context as a whole may conspire to generate a more detailed, specific characterization, and thereby a richer metaphor. In particular, speakers often employ a single, central verb metaphorically, and then extend, specify, and enrich that metaphor by employing other terms which are linked to the central verb, either by colligation or simply through our more general knowledge about the relevant sort of situation.[14] Even verbs with

---

[13] Cf. Searle, "Intentionality, Consciousness and the Background."

[14] Cf. Lynne Tirrell, "Extending: The Structure of Metaphor." See also Kittay, *Metaphor,* especially chapter 7. Another example of an implicit aspect introduced by colligation over an extended metaphor is the third stanza of T. S. Eliot's "Love Song of J. Alfred Prufrock":

> The yellow fog that rubs its back upon the window-panes,
> The yellow smoke that rubs its muzzle on the window-panes,
> Licked its tongue into the corners of the evening,
> Lingered upon the pools that stand in drains,
> Let fall upon its back the soot that falls from chimneys,
> Slipped by the terrace, made a sudden leap,
> And seeing that it was a soft October night,
> Curled once about the house, and fell asleep.

The entire passage is unified by the unspoken aspect of an animal, most likely a cat (though perhaps a dog) being applied to the yellow fog. The governing characterization is much more specific if we treat the passage as a whole rather than on a sentence-by-sentence basis: each sentence contributes different features to the governing

broad applicability can be extended in this way. In (17), just one relation, between the chair and the discussion, was introduced by colligation, and nothing more was specified for the governing characterization about the particular sort of agent doing the plowing or the particular sort of object being acted upon. By contrast, in Anthony's famous speech to the Romans in *Julius Caesar*,

(18) You all do know this mantle...
    Look, in this place ran Cassius' dagger through;
    See what a rent the envious Casca made;
    Through this the well-beloved Brutus stabb'd,
    And as he pluck'd his cursed steel away:
    Mark how the blood of Caesar followed it,
    As rushing out of doors, to be resolv'd
    if Brutus so unkindly knock'd or no;
    For Brutus, as you know, was Caesar's angel.[15]

we have a much more complex collection of relations, all introduced by way of a single governing characterization. By themselves, 'rushing' and 'following' can apply literally to very many, very different sorts of thing in rather different ways, including blood, and so do neither generate highly specific, nuanced characterizations nor trigger aspectual thought.[16] The addition of the phrase 'out of doors' introduces the ancient and familiar metaphor of the body as a house, which itself suggests many subsidiary matches. But we still lack a specific characterization for the motion of the blood itself, which is the central subject of the central metaphor. Once Shakespeare introduces the final phrase, however, we have a much more specific governing characterization for thinking about of the blood and its motion. The blood is being claimed to rush in a way that is to be understood under the aspect of a resident of the house who chases after someone who has rudely interrupted the quiet with a loud, unexpected knock. Brutus's stab and the consequent flow of blood are to be thought of together under the single overarching aspect of this situation.

---

characterization, which can then be applied to interpreting the metaphors in the other sentences. Stern uses this example to make a similar point about what he calls 'enthymematic' metaphors (that is, metaphors which don't make explicit their essential guiding presuppositions) in his discussion of extended metaphors; *Metaphor in Context*, p. 174.

[15] Shakespeare, *Julius Caesar*, III.ii.170ff.

[16] I take it that blood literally does rush out of bodies on certain unfortunate occasions, and that 'follow' literally means just something like 'come after'.

I mentioned in §2.1 that this case introduces a puzzle about what is to be seen as what. John Crowe Ransom suggests that the blood is being described in the terms of a page rushing out of doors, while Monroe Beardsley excoriates Ransom for introducing unwarranted imagery into the metaphor.[17] I think we can now see that in a sense, both are correct. Beardsley is correct that there is no specific *conventional* association between pages and rushing out of doors, as there was with ballet dancers and pirouetting. The phrase 'rushing out of doors' merely requires, as a matter of semantics, that someone or something do the rushing. Once we include the later phrase about 'being resolved', then we know that the governing characterization must be of a *person* rushing, because only people can engage in mental activities like resolving or ascertaining. Still, no semantic fact necessitates or justifies Ransom's introduction of the page specifically into the situation characterized by the governing characterization.

Nevertheless, Ransom does have good reason for introducing the page. If we cultivate the aspect of the relationship between a person and a house as a means for thinking about the relationship between the blood and the body, it is appropriate to match the blood's role in the body to the role of a servant in the house, and to treat the mind as the body's master. (This is perhaps especially apt given that Caesar is an emperor.) And if we think further about the specific situation of a resident of the house chasing after someone who has rudely knocked, then a servant seems to be the most likely candidate for doing the chasing. Armed with a characterization of how a page might rush out of doors to investigate the rude disturbance of his master's serene privacy, we achieve a substantially richer and more specific understanding of the subject characterization than if we simply stuck with the more general governing characterization that is explicitly warranted by the semantics of 'rushing out of doors'. So, on the one hand, by the strictest standards Ransom does introduce unnecessary imagery into his interpretation. But on the other hand, by standards which are sensitive to the aesthetic interests and the communicative

---

[17] Ransom writes: "The blood flowing from the blood-vessels which have been punctured becomes the page or houseboy rushing out of Caesar's body, the house, in order to make certain that this is really Brutus who is making such a rude disturbance to gain entry where he is always welcome" ("Poetry: I. The Formal Analysis," p. 452). Beardsley responds: "Now there is obviously no page in these lines, any more than there is a rudely awakened householder or soon-to-be embattled farmer alarmed by Paul Revere." ("The Metaphorical Twist," p. 265).

purposes of the speaker in the passage, Ransom is correct that a page is a quite fitting candidate for doing the rushing, and so his introduction of the page into the governing characterization is warranted.

One lesson of this discussion about verb colligation and extended metaphor, then, is that it is potentially both beneficial and dangerous. The great benefit of introducing a governing characterization indirectly by verb phrase colligation is one can thereby compress a lot of information about the *subject* being described into the metaphorical specification of the *action* or state which the subject is undertaking or undergoing. The metaphor is thus that much more efficient. But the danger of introducing a governing characterization indirectly by verb phrase colligation is that unless the association between the action (or state) and the agent or patient is highly conventional and specific, like the association between 'pirouette' and ballet dancers, then the resulting governing characterization may not be that specific. This leaves open the possibility that in the attempt to make it more specific, and thereby to generate a more powerful metaphorical interpretation, hearers may end up misinterpreting the metaphor by introducing into the governing characterization a constituent other than the one the speaker intended.

5.2.4: Whole-Sentence Metaphors

In my discussion of metaphorical utterances which fail to fit the simple case, I have so far introduced three sorts of complexity. **First**, the sentence uttered may not fit the simple syntactic model of '*a* is *F*': it may include quantified phrases, including definite descriptions. I said that the process of metaphorical interpretation still proceeds in the same basic manner in these cases. **Second**, the noun phrase may itself be metaphorical; in this case the subject under discussion, and so the relevant subject characterization, must be identified indirectly, partly by reference to what the speaker might be talking about at this point in the conversation, and partly by reference to what could appropriately be described in such metaphorical terms. **Third**, the individual or type of individual which the *governing* characterization characterizes may be introduced indirectly, by colligation from the verb phrase.

The **fourth** sort of case that I want to discuss combines the complexities introduced by the last two. These are 'whole sentence metaphors', in which the complete sentence uttered, construed literally,

specifies a whole situation which is to be characterized by the *governing* characterization. The whole situation which is to be characterized by the *subject* characterization, on the other hand, must be identified indirectly in the manner described above for metaphorical noun phrases.

Consider, for example,

(19) The rock is becoming brittle with age.

This sentence might be uttered among a group of students as they emerge from meeting with an aging but eminent professor.[18] In a context like this, the professor might be just one salient object among several. But the relatively familiar metaphorical use of 'rock' to mean a person of central and foundational significance (cf. Jesus' description of Peter as "the rock upon which I build my Church"), and the obvious applicability to the professor of the prominent rock-features of solidity and durability, combine to make the professor the most likely subject of the utterance. The basic aspect, then, characterizes the professor as a rock. However, the utterance makes a more specific claim: not just that the professor is a rock, but that the professor is becoming *something* with age, where the relevant quality is to be identified by applying the governing characterization of a rock becoming *brittle* with age. (Notice that 'with age' is included in both characterizations.) Given the basic aspect of the professor as a rock, 'brittle' then introduces by analogical matching the feature of his being intellectually inflexible (and perhaps personally irritable). As with (13), we might express the speaker's intended meaning in appositive form:

(19') The professor, a person of foundational significance, is becoming intellectually inflexible with age.

Similar analyses apply to other whole-sentence metaphors, such as

(20) The fox has slipped its snare,

said of George who has escaped the embarrassing intellectual comeuppance his colleagues had set for him, or

(21) The sun is blazing,

said of Achilles as he rages upon the battlefield, or

---

[18] This sentence was initially offered by Michael Reddy ("A Semantic Approach to Metaphor") as an example of a sentence with no semantic deviance: it could be uttered literally and perfectly felicitously during a geology expedition.

(22) Hark! The sun's first light breaks through yonder window,

said by Romeo in lieu of the famous beginning to his speech.

Although it is not exactly a unified whole-sentence metaphor, I think

(23) Confusion now hath made his masterpiece,

uttered by Macduff in Act II, scene iii of *Macbeth*, works in this same way. The governing

characterization, introduced in part by colligation, is of *an artist who has now made his masterpiece.*

Although 'confusion' is intended literally, it refers to such an abstract notion that it doesn't provide us

with much guidance in generating the relevant subject characterization; all we can provide initially is

confusion, which has now done something. It is contextually salient, however, that King Duncan's

murder has just been discovered. This suggests that the appropriate subject characterization is *confusion,*

which has now created this regicide. A more transparent way to reveal the aspect's application might

then reinterpret the metaphor into two appositional phrases:

(23') Confusion, who is an artist, now has made his masterpiece, which is this regicide.

We are supposed to think of the regicide as an artwork, the product of a devilish intentional design which

will produce yet more confusion. (Thus, confusion is in a sense both the agent and the product of the

regicide.)[19] At the same time, we are supposed to think of confusion as an agent at work in the affairs of

men. The characterization of men's actions (specifically, the murder) as the product of another, higher

agent in turn makes those men seem more like the objects than like agents in their own right.

Before concluding this section, I want briefly to take up one final class of cases, which I think

represent more of a curiosity than a genuine challenge. These are what we might call **'tautologous**

metaphors', such as

(24) War is war.
(25) Boys will be boys.

(Intuitions appear to be mixed whether these are really metaphors at all.) The puzzle is that because the

same term appears in both the noun and verb phrases, there seems to be no way to supply distinct

---

governing and subject characterizations, and so no way to generate an aspect.[20] I think what happens here is that the hearer is intended to apply the *stereotypical* characterization of the class described as an aspect to the hearer's own characterization. In effect, the speaker claims that the class described really does possess the features that are most prominent in our shared stereotypical characterization of it.[21]

In summary, there are two main ways in which the simple relationship I described in §5.1, between the sentence actually uttered and the appropriate aspect, must be extended to cover more complex cases. In the basic case, the sentence's noun phrase denotes the subject of the subject characterization in literal terms, while the verb phrase directly and literally indicates the governing characterization. In more complex cases, either the subject or the governing characterization, or both, is specified only indirectly. In such cases, the hearer must supply those characterizations by relying on either the contextual salience of objects and events, or the usual colligates of the words uttered, or both.

5.3: *From Aspect to Content: Complexities*

I've now surveyed all the ways in which a sentence intended metaphorically can generate the appropriate aspect. The relationship between the aspect and the speaker's intended content also introduces important variations among metaphors. I think that insufficient attention to these variables generates much of the disagreement over metaphor: different theorists take different metaphors as paradigmatic, overgeneralize from the limited cases they consider, and then conclude that other theorists, who are driven by intuitions gleaned from other metaphors, must be just wrong. I hope that by making these variables explicit, I can show that a basic unity underlies them after all. This will stand us in especially good stead when I take up the question of whether metaphors can be paraphrased, in §6.1.

---

[20] Thanks to David Davies for raising this question.

[21] I encountered another metaphor of this form in the title to the book *No Island is An Island: The Ordeal of Martha's Vineyard,* by Anne W. Simon. Although interpretation here seems in fact to proceed by way of the sentence "No man is an island," if we suppose that Donne's metaphor were not itself already so familiar, then the basic logic of the new metaphor would seem to be like those above. It denies that islands (and one island in particular) have the properties which are most prominent in our stereotypical characterization of them: isolation, self-sufficiency, and so on.

5.3.1: *Conversational Weight*

The first, and perhaps most obvious, difference among metaphorical utterances is the *conversational weight* they are accorded in their context of utterance; I already mentioned this variable in §5.1. Some metaphors are tossed off by-the-by, in the course of making another point; some are employed in the give-and-take of conversation to make a point of their own, but with an eye to the central conversational topic. Other metaphors are offered as an important point to be explored further; and still others, like Romeo's utterance of (1), become the center of the conversation itself. The greater the conversational weight borne by a metaphorical utterance, the more effort the hearer is expected to put into interpreting it. Greater expected effort raises the number of features the hearer is expected to uncover, by lowering the threshold for what counts as 'the *a*-features *most* tightly matched to the *most* prominent *F*-features'. Whereas the speaker commits herself to the applicability of only a few such features for 'incidental' by-the-by metaphorical utterances, greater conversational weight results in a richer content for the speaker's meaning.

One consequence of this conversational dynamic is that speakers usually employ as very lightly-weighted utterances only metaphors in which a very few, highly prominent *F*-features obviously and tightly match a few *a*-features. The conversational context does not warrant much interpretive effort, and so the relevant feature must be easily recoverable. A good example of this is, I think,

(26) The Dean's house is a Taj Mahal,

used to mean that the Dean's house is luxurious and rather opulent. In these cases, the conditions for the truth of what the speaker meant are the most clear-cut.

By contrast, the speaker's meaning in heavily-weighted metaphorical utterances tends to be more open-ended and indeterminate. In some contexts in which the metaphor itself becomes the center of the conversation — paradigmatically, in poetry — the speaker intentionally cedes some of the responsibility for determining that content to the hearer. She assumes that the hearer is a cooperative partner in the project of developing the metaphor's content, and so that the hearer has taken her characterizations on board as fully as he can. Given this assumed shared background, the speaker allows that whichever

218

subject-features the hearer finds to be appropriately matched to features in the governing characterization will count as part of the metaphor's meaning. She asserts that her utterance will be true when given that meaning. The hearer may thus uncover legitimate *M*-features which the speaker didn't expect, and perhaps even would have rejected.[22] The fact that the hearer is still constrained by the speaker's characterizations, though, means that there remains an important difference between even this quite open-ended and participatory interpretation and the cultivation of an aspect for one's own intellectual and aesthetic interest. Thus, while engaging with poetry does call for a 'free play of the imagination', this 'play' still occurs within the ambit of the speaker's intention.

5.3.2: *Independent Expressibility*

The next important variable among metaphors' contents is how easily the intended *M*-features are to express by literal linguistic means. To the extent that they are expressible in this way, the metaphor itself is 'ornamental'. Thus, for instance, if I utter

(27) Richard is a lion

in a conversational context in which the utterance doesn't carry very much weight, I may mean no more than that Richard is courageous (though I may also mean to ascribe a specifically lionish sort of courage to him). Because our standing vocabulary includes the word 'courageous', the metaphor is relatively disposable: I might as well have said explicitly and literally what I meant.

Sometimes, however, there is no conventional linguistic alternative. In many such cases, demonstrative reference can provide literal access to the intended feature. However, it can also happen that the speaker has in mind a property *P* with which she has had the sort of causal interaction necessary to underwrite a demonstrative thought about *P*, but that *P* itself is not instantiated in the immediate context of utterance. In such a case, if her hearer has not also had the right sort of interaction with *P* — or even just if the speaker and hearer lack a shared method for identifying *P* — the speaker may still be able

---

[22] David Hills suggests (in conversation) that in some such cases, the speaker engages in a shallow pretense of intending some determinate content in order to solicit from the hearer some content that's *worth* intending.

to refer to *P* demonstratively, but her hearer will be unable to understand what she is talking about. There will thus be no literal means by which the speaker can communicate her intended content.

Further, as we saw in §3.5.1, demonstrative reference — even demonstrative reference supplemented by a general term, such as 'that color' or 'this sort of wintry day' — does not always make clear just what it is about the object or feature denoted that the speaker wants to point out. Indeed, the speaker herself may not be entirely clear about just what this is. She may have a genuine ability to identify that the relevant feature *is* instantiated when she is confronted by it, but not be clear about just how to isolate that feature from within a more complex pattern. In such a case, her ability to identify the relevant feature is still essentially anchored to the original object that exemplifies it, but she now wants to predicate that feature of some other thing.

In this situation, one convenient way — indeed, in some cases, the only way — to indicate to her hearers what the relevant feature is may be via the role which an analogous feature plays in a shared characterization of some *other* object. As examples, think of trying to describe the movements of a ballerina, or the feelings aroused by a piece of music. In the first case, we lack the physical skills to mimic the movements in question; in the latter, we cannot ostend our emotions. However, in such a situation, a metaphor like

(28) She was a swan skimming over a smooth lake

or

(29) The Fauré Requiem puts you on the top of a lonely, icy mountain crag

may still succeed in communicating what the property in question is.

In these cases, at least as I am imagining them, the speaker wants to specify a particular feature, but lacks the literal linguistic resources to communicate this feature to her interlocutor. Perhaps more commonly, one might also have a particular *constellation* of features in mind, but lack the cognitive capacity (or simply not want to put in the required effort) to identify each of those features individually.

One thinks of them as an "undifferentiated chunk" rather than as a complex conjunction.[23] A metaphor, even a relatively unpoetic one like

(30) Steve is a sheepdog,[24]

may succeed quite nicely in communicating such a 'chunk' of properties. In this case, the speaker is after the particular combination of loyalty, adoration, fierceness, subservience, enthusiasm, and perhaps a certain sort of posture and look in the eyes, though she hasn't isolated each of these features in this explicit way. Because characterizations are precisely constellations of such multiplicities of features, they provide useful vehicles for communicating 'chunks' of features when the speaker has not explicitly reflected upon just what the constituent elements are.

Finally, we may be interested in identifying an as-yet unnamed feature which plays some as-yet only hazily identified role in a complex field of features. We can't directly ostend the feature we're interested in, because it's accessible to us only through the causal (or other) role it plays in the complex field. But we're not yet sufficiently clear on just what this causal role is to be able to identify the feature definitionally. Metaphor may then provide our only means for referring to such a feature. Richard Boyd argues that theorists investigating fundamental cognitive processes are in this position.[25] We have reason to believe that *some* feature plays a role in consciousness which is analogous to a feedback loop in a computer program, or that *something* plays a role in cognitive data processing analogous to folders and files. But we don't yet know enough else about these features to investigate them directly and on their own, or even to isolate in just which respects that cognitive feature will be analogous to the computer model. Thus, the metaphor plays an essential role in fixing what we are talking about.

---

[23] The phrase is Stern's (*Metaphor in Context*, pp. 191-192).

[24] The example is from Patti Nogales, *Metaphorically Speaking*, p. 1.

[25] "Metaphor and Theory Change: What is 'Metaphor' a Metaphor For?," pp. 367-369. As Thomas Kuhn emphasizes, the need for metaphor is especially pressing when we want to refer to natural kinds rather than to individuals, because in that case "one loses access to the career line or lifeline which, in the case of proper names, enables one to check the correctness of different applications" ("Metaphor in Science," p. 411). That is, one act of ostending a kind or a quality does not establish clear identity conditions in the same way that a single act of ostension intuitively does for an individual. For more on this difference, see Scott Soames, *Beyond Rigidity: The Unfinished Business of Naming and Necessity*, chapters 9 through 11.

5.3.3: *Aspects' Distance, Resonance, Novelty*

These first two sorts of variation among metaphors have concerned the speaker's intended content: how much or how little goes into that content, and how easily that content can be expressed in other terms. The three variables I discuss in this section concern the aspect itself. First, we have an aspect's *distance* from our normal characterization of the subject. An aspect *F* is more distant from the subject *a*'s normal characterization to the extent that *a* and *F*-type things are of very different sorts. Thus, *being the sun* is a more distant aspect as applied to Juliet than *being a queen* is. Second, an aspect is more *resonant* to the extent that it generates many matches. Third, an aspect is more *novel* to the extent that the class of *F*-type things is not usually used as an aspect for thinking about *a*-type things. (The relevant class here will vary considerably, but it will typically focus on conventionalized contrast classes: for example, the four seasons, the colors, the ages of man, familiar animals, or sorts of machines.)

Often, all three of these variables go together — especially for what we think of as good metaphors. Thus, for instance, an utterance of

(31) A geometrical proof is a mousetrap

is striking because the aspect is so distant from the subject, but it is still relatively resonant and novel. However, none of these variables is essentially correlated with the others. Some metaphors with distant aspects are relatively well-established, such as the general aspect of celestial bodies — the sun, the moon, the stars — applied to people. And some novel metaphors, even rather good ones, are not particularly resonant. Goatly cites

(32) The seas tented up his oilskin,

used to describe the way a fisherman's raincoat was swelled up by the air trapped under it as he sank deeper into the sea, as an example here.[26]

This example brings out a particular way in which a metaphor may be powerful, effective, and even novel, without being resonant. We don't have a very rich, nuanced, detailed characterization for tents, although we do have some familiarity with them, and in particular with the way they look. (In other

---

[26] *The Language of Metaphors*, p. 97.

cases, we may have a rich, nuanced governing characterization, but the aspect may be so distant that very few matches are available between the two things.) Specifically, I doubt that the characterization of tents which we do have captures much in the way of intuitions about centrality: it is true that some features of tents cause other of their features, but this is not something we have strong commitments about the way we do about people's personalities. Nonetheless, 'tent' can still be used as a metaphor (in addition, it can be turned into a verb, a common though quite distinct form of non-literal word usage). Here, it is just the match between the respective shapes of a tent and the fisherman's coat which is exploited. Because we lack a standing literal term for the creation of this shape, 'tented' is an effective metaphor: a compact means for communicating a claim which would otherwise be difficult to state. More generally, a few very specific matches, especially involving experiential features, may contribute to a metaphor's effectiveness at least as much as a nuanced, rich structure of matches can.

If we now consider what it means for a metaphor's aspect to be **novel**, we in turn need to distinguish between two sorts of familiarity. **First**, there are what M. B. Dagut calls "embalmed metaphors of literature."[27] These are metaphors which begin as the novel creations of a single author, and achieve familiarity because the author's work becomes a classic, but whose generating aspects do not enter the common repertoire of everyday speech. Shakespeare's metaphor of 'salad days' might be one example here; another might be the characterization of one's lover as a garden, employed in *Song of Songs* and in love poems from many different cultures; a third might be the metaphor of men's lives as leaves on a single great tree, exploited by Homer, Virgil, and Dante. We don't normally and naturally employ metaphors relying on these aspects in everyday speech. **Second**, there are aspects whose applications have become so 'routinized', whatever their original source, that we do use them without reflection in ordinary conversation. The description of people as different sorts of animals — pigs, lions, roosters, bears, gorillas — are paradigmatic cases here. 'Embalmed' metaphors may be quite familiar without ever becoming 'routinized', even if they become established tropes or even clichés within the literary tradition.

---

[27] "Can 'Metaphor' be Translated?," p. 23.

Notice, further, that the novelty of a generating aspect is not necessarily commensurate with the novelty of a metaphor which exploits it. A surprisingly large proportion of poetic metaphors, perhaps especially Shakespeare's, employ well-entrenched aspects in novel ways and to novel ends. We encountered this in §5.2.3, where I cited Shakespeare's exploitation and extension of the well-entrenched aspect of houses for thinking about bodies, in order to characterize Brutus's stab and the consequent flowing blood in a novel way. Sonnet 73 is a stunning, yet more extended example of this sort of 're-vivification':

> That time of year thou mayst in me behold
> When yellow leaves, or none, or few do hang
> Upon those boughs which shake against the cold
> Bare ruined choirs where late the sweet birds sang.
>
> In me thou seest the twilight of such day
> As after sunset fadeth in the west,
> Which by and by black night doth steal away,
> Death's second self, which seals up all in rest.
>
> In me thou seest the glowing of such fire
> That on the ashes of his youth doth lie,
> As the deathbed whereon it must expire,
> Consumed with that which it was nourished by.
>
> This thou perceiv'st, which makes thy love more strong,
> To love that well which thou must leave ere long.

Here Shakespeare employs metaphors based on three entrenched, even cliched aspects for thinking about a lifetime: in terms of the seasons of the year, in terms of the hours of the day, and in terms of the stages of a fire. But he employs novel metaphorical expressions and exploits novel specific matches for those familiar aspects. Thus, he describes autumn as a time when nearly leafless boughs "shake against the cold," as humans rail against old age. And he describes one's death-bed as a place where one lies "on the ashes of his youth...consumed with that which it was nourished by," as our aged dissipation is the product of our youthful enthusiasm and excess. Further, by juxtaposing and intertwining those aspects, he reciprocally leads us to think of the seasons, the day, and fire themselves in new ways as well.

5.3.4: *Commitment to the Aspect*

The final pair of variables I want to take up are perhaps the least considered in discussions about metaphor; the second is, for me, one of metaphor's most intriguing facets. I have already discussed both variables in the context of aspectual thought *per se* in chapter 4. First, recall from §4.5.1 that an aspect is *external* when the thinker cannot bring the concept of *being F*, which is associated with the relevant governing characterization *F*, within the subject characterization as fitting for that subject. In such a case, the thinker simply has no intuitions about how the subject would be if it were in fact *F*, let alone how it might be *fitting* for the subject to be *F*. *Being the sun* is, I take it, an external aspect for Juliet, whereas *being a queen* is an internal aspect for her.

Second, recall from §4.7.3 that whether the aspect is internal or external, there are two sorts of commitment one might have to it. On the one hand, one might be committed to the claim that one's overall characterization of *a* as it is structured under the aspect *F* accurately reflects *a*'s true character. On the other hand, one might merely be committed to the aspect's instrumental utility. Although it suggests a distorted overall picture, one still finds it useful because it brings out certain otherwise neglected, perhaps even inaccessible, features of *a*.

When these differences among aspects are translated into the realm of metaphor, they introduce important complexities in what a speaker can commit herself to by speaking metaphorically, and even in what counts as a metaphor at all. I have by now discussed three levels of what might loosely be called **'content' in relation to metaphor. First**, there is semantic content — the content of what the speaker actually says — which is determined by the conventional meanings and mode of combination of the words she utters.[28] **Second**, there is the characterization of the subject, with the particular structure that is imposed on it through the application of the relevant aspect. This is the level of 'content' about which two thinkers disagree when, say, one dismisses some feature *P* of the subject's as a mere fluke, while the other insists that it is central and thus revealing of the subject's true character. **Third**, there is pragmatic

---

[28] Again, where this includes contextual factors made relevant by the conventional meanings of constituents in the sentence uttered.

content — the content to which the speaker primarily intends to commit herself by means of her utterance.

On my view, a speaker may be committed to — and make a speech act which has as its content — one, two, or even all three of these levels of 'content', while still speaking metaphorically. The possibility of such multi-leveled commitment contradicts the usual, loose definition of metaphor as a speech act in which a speaker necessarily says something she *doesn't* mean in order to mean something else, where the connection between what's said and what's meant is mediated by some vague sort of resemblance relation.

This traditional definition takes as its paradigm a metaphor like

(33) Sam is a pig,

and treats metaphor exclusively as a figure of *speech*. In my terms, in a case like this the speaker is committing herself only to the content given by predicating of *a* the *M*-features determined by the aspect. Here, both the sentence uttered and the aspect it generates are merely tools for getting the hearer to isolate the relevant *M*-features, and thereby to identify the speaker's intended speech act. The speaker certainly doesn't believe, not does she want to be taken to believe, that Sam really is a pig, nor even that his whole character is piggish. Perhaps she merely means to claim that he is a glutton.

As we saw in §1.1.1, many traditional theories of metaphor assumed that all metaphors are like (33) in being categorically or at least absurdly false. And many metaphors do in fact work in this way. However, recall that Ted Cohen and others have shown that metaphors can also be 'twice-true': true when construed literally, and also true when construed metaphorically. Putative examples are

(34) No man is an island
(35) George W. Bush is a primate.[29]
(36) Moscow is a cold city.

(Whether one counts these as examples of 'twice-true' metaphors depends on one's views about the truth-values of the relevant metaphorical claims.) In these cases, speakers are in a weak sense committed *both* to what's said and to what's meant: they believe that both propositions are true. However, what is

---

[29] John Searle suggested this example to me.

226

literally said by utterances of (34) or (35) is so trivially true that no one is likely to intend to communicate

their contents; they are unlikely to provide a relevant contribution to a conversation. Thus, these cases are

still like the traditional paradigmatic case of (33) in that the sentence uttered, construed literally, plays a

merely instrumental role. It is not part of what the speaker intends to commit herself to by her utterance,

because she takes it already to be part of the common conversational ground. By contrast, an utterance of

(36) *might* be intended to communicate the content it literally expresses. David Hills calls cases like these

'twice-apt',[30] because both what is said and what is meant may be appropriate in the conversational

context.[31]

But now, having distinguished the three levels of 'content', we can push this point about 'twice-

over' metaphors further. I take it that Romeo's utterance of

    (1) Juliet is the sun

is also 'twice-apt' or 'twice-meant', but in a rather different sense. Romeo is committed *both* to the

appropriateness of the aspect generated by his utterance and to the truth of the speaker's meaning it fixes.

He thinks both that the aspect generates an appropriate overall characterization of his beloved, and that it

enables him to make a true claim about her. I think, further, that Romeo is committed to *claiming* that

Juliet has the character revealed by the aspect: that is, that the overall characterization of her under the

aspect is true. That is, if someone merely agreed with Romeo that Juliet is beautiful, good, and nurturing,

but disagreed that these were prominent and central features of who she is, then I don't think Romeo

would accept that they were in agreement.

Romeo is clearly not, however, committed to the truth of what he actually *says*: that Juliet is the

sun. If he were committed to this claim, then he ought to be institutionalized for grand dementia and

delusion.[32] Given these two different sorts of 'twice-apt' metaphors, though, we might now wonder

---

[30] "Aptness and Truth in Verbal Metaphor," p. 147.

[31] Consider also an utterance of 'Those were my salad days,' intended to claim both that those were my young, more hippy-liberal days in which I ate a lot of salad, and also that those were the days in which I was myself fresh, näive, and energetic. Here the two contents are very different, but both are intended.

[32] Shakespeare has Hamlet play with just this ambiguity at several points in *Hamlet*, including especially in his dialogue with Polonius, at II.ii.194ff, which begins "Words, words, words...'.

whether there could be other cases of metaphor which would count as '*thrice*-apt'. In such a case, there would be no 'deviance' in the utterance at any level. In particular, while with (1) the aspect is external, 'thrice-apt' metaphors would have to generate internal aspects. The speaker would be claiming that *a* really is *F*, that *a*'s overall character is accurately reflected by thinking of *a* under *F*'s aspect, *and* that *a* is $M_1, M_2, ...M_n$, where these features are identified in the usual way. Three putative examples of such metaphors are

(37) Jesus was a carpenter;[33]
(38) He was a man, all and all,

said by Hamlet to mean that his father was all that a man can be, a true man; and Marlowe's claim, in Conrad's *Heart of Darkness*, that

(39) The French man-of-war fired into the continent, shelling the bush,

where we are to understand the ship as both literally *and metaphorically* attacking Africa.[34] I can also almost imagine a context in which (35) above might be intended in this way, as a 'thrice-apt' utterance. The speaker would then not just claim that George W. Bush was a wild, unrefined beast, as he would if he intended the aspect merely instrumentally. Rather, he would also claim that Bush's entire character is clearly revealed by thinking of him as a primate: his attitudes toward defending and acquiring territory, his displays of anger in defending territory, his patterns of aggression and reconciliation, his social grooming behavior, his personal appearance, and so on.

Of these four cases, I find (37) to be the most convincing. Here there is enough of a difference between what the speaker says and what she means that the latter seems genuinely metaphorical. (38) is more like the trivially twice-true (34): no one is likely to want to claim that Hamlet's father was a man (though we can imagine that someone might claim that he *is* a man — rather than a ghost). In this respect, it is much like the 'tautologous metaphors', such as

(24) War is war,

---

[33] Cohen, "Notes on Metaphor," p. 254.
[34] The second two examples are from Hills, "Aptness and Truth in Verbal Metaphor," pp. 130-131. The Hamlet example traces back to Empson's discussion of 'pregnant' uses of words.

which I discussed in §5.2.4. Given the context I am imagining for (35), it seems to be simply a literal statement which, while not itself informative, is uttered with the intention of inducing the relevant aspect in the hearer, in much the same way that I said we might tell a revealing story. Thus, it is not clear that the speaker really is intending to claim what is literally said by an utterance of (35). Finally, my intuitive response to both (38) and (39) is that they are cases of **symbolism** — that is, very roughly, as cases in which a particular concrete fact or quality both instantiates and represents a larger, more complex and abstract situation.

If this is right, then perhaps we can analyze (37) in this same manner: perhaps Jesus' being a carpenter symbolizes his relation to mankind. If we decide this is not an instance of symbolism, but is a genuine metaphor after all, then we may have to revise my initial definition of metaphor: as an utterance in which the speaker does not *primarily* intend to communicate what she says, and where the connection between what is said and what is meant is mediated by aspectual thought. (The insistence on 'primarily' here is intended to account for trivial 'twice-true' metaphors like Cohen's.) However, it seems more likely that the line between metaphor and symbolism here is fuzzy, marked by how different the literal and the 'other' meaning are from one another. A more definite resolution of this matter would obviously require a significantly more developed theory of symbolism.

If we return from the possibility of 'thrice-apt' metaphors to mere 'twice-apt' metaphors like Romeo's utterance of (1), I think they point the way to explaining an odd but important and revealing fact about our response to some metaphors. This is our inclination with some metaphors to say that we mean *just* what we say, even as we acknowledge that the utterance is literally false. In these cases, I think, the speaker is committed to the accuracy of the overall characterization imposed by a metaphor's generating aspect. She therefore wants the hearer to take the metaphor seriously, as the key to a more general understanding of the subject's true character, and not just as an instrumental means to identify some distinct content.

One provocative way to put this point is that the speaker wants to take the metaphor as *literally* as possible — even while acknowledging its literal falsity. Samuel Levin describes this phenomenon by

Here is an example of that phenomenon. Suppose you and I are hiking in the hills, and come across a lake, whose shore is filled with springtime flowers. I utter

(40) The daffodils are dancing in the breeze.

My utterance is obviously literally false — flowers have no legs, and do not dance — but it's clear what I mean. The flowers are moving in a way that, like dancing, is quick and rhythmic. The breeze, the daffodils, and their motion, are all seen by both of us in their usual ways. I have chosen a slightly colorful but quite efficient means of bringing the salient features to your attention.

Now, however, suppose that you read Wordsworth's "I Wandered Lonely as a Cloud":

> I wandered lonely as a cloud
> That floats on high o'er vales and hills,
> When all at once I saw a crowd,
> A host, of golden daffodils;
> Beside the lake, beneath the trees,
> Fluttering and dancing in the breeze.
>
> Continuous as the stars that shine
> And twinkle on the milky way,
> They stretched in never-ending line
> Along the margin of a bay:
> Ten thousand saw I at a glance,
> Tossing their heads in sprightly dance.
>
> The waves beside them danced; but they
> Out-did the sparkling waves in glee:
> A poet could not but be gay,
> In such jocund company:
> I gazed — and gazed — but little thought
> What wealth the show to me had brought:
>
> For oft, when on my couch I lie,
> In vacant or in pensive mood,
> They flash upon that inward eye
> Which is the bliss of solitude;

---

[35] "Standard Approaches to Metaphor and a Proposal for Literary Metaphor"; *Metaphoric Worlds: Conceptions of a Romantic Nature*.

> And then my heart with pleasure fills,
> And dances with the daffodils.

Nearly the same sentence as I uttered appears again within the poem; but in the new context, the same metaphor is much more deeply meant. That is, Wordsworth intends for us to cultivate the generating aspect itself as fully as possible: to seek as many matches that are as tight and direct as possible, even if this requires transforming our characterization of daffodils quite significantly. In particular, we are expected to make matches between features of dancing people and of daffodils that we would have discarded as outlandish in the previous context. It's normally obvious that daffodils aren't energetic and full of enthusiasm, that they don't take pleasure in social interactions with other daffodils. But as we dwell on the poem, we begin to think of the daffodils as personified, and it begins to make sense to think, say, that the daffodils might compete together against the lake's waves to be most lively and gleeful.

Of course, even after this imaginative conceptual transformation, many features of dancers and daffodils remain unmatched. We do not think that daffodils have legs made up of bone and muscle, and that they dance by contracting those muscles, for instance. Still, one *might* come to see the daffodils as analogous to human bodies, not merely in their style of movement, but in being filled with 'the sap of life' just as humans are filled with blood, or in standing on the earth with a certain jaunty stance just as humans do before a dance. And as we cultivate these previously outlandish matches, we also begin to think of nature more generally as not just alive but as full of emotion and intention. In short, we begin to adopt the vision of Romantic naturalism championed by Wordsworth. As Samuel Levin says, "on this view [of poetic metaphor] it is not the language that is remarkable, it is the conception; given the conception, the linguistic description of it follows straightforwardly."[36]

5.3.5: *'Poetic', 'Ordinary', and 'Dead' Metaphors*

Thus far in this section, I have laid out **seven variables** according to which metaphors can differ:

(a) the weight they carry in the conversational context;
(b) the *M*-features' independent semantic expressibility;

---

[36] "Standard Approaches to Metaphor and a Proposal for Literary Metaphor," p. 133.

(c) the distance of the aspect from our normal characterization of the subject;
(d) the resonance or fruitfulness of the aspect: how many matches it supports;
(e) the novelty of the aspect for thinking about the subject;
(f) the aspect's being external or internal;
(g) the sort of commitment that the speaker has to the aspect itself.

Typically in discussions of metaphor, these variables are collapsed into a single dimension of difference: metaphors can be 'ordinary' and conversational, they can be 'rich' and poetic, or they can perhaps fall somewhere in between.

There is something to the intuition of this single sort of difference, because many of these variables tend to be correlated. At the one extreme, 'ordinary' metaphors carry little conversational weight, employ a familiar and relatively non-resonant but external aspect in a merely instrumental fashion, and use that aspect only to isolate a few features which could be described by standing terms in the language.

(27) Richard is a lion.
(33) Sam is a pig.

are prime examples of this sort of metaphor. At the other extreme, 'poetic' metaphors occupy the center of conversational attention, in which the hearer and speaker jointly explore the rich implications of a resonant and relatively novel aspect, and the speaker is committed to the claim that the generating aspect accurately represents the character of the subject under discussion.

(1) Juliet is the sun

is a paradigmatic example here.

Once we've distinguished these seven features, though, we can see that metaphors can be more or less 'ordinary' or 'poetic' in different ways. A metaphorical utterance may carry little conversational weight, but use a distant and novel but non-resonant aspect to pick out an relatively inaccessible feature; I cited

(32) The seas tented up his oilskin

as an example of this. Or a metaphorical utterance may occur as the guiding metaphor in a poem, say, but rely upon a familiar aspect to which the speaker is committed only instrumentally, in order to derive a

limited number of features which could be stated in plain words. (I suspect that many bad poems work in this way, but I haven't yet encountered a particularly satisfying example.) Often the shorthand of 'ordinary' and 'poetic' is indeed useful and causes no real harm, but we need to remember that it is just a shorthand. Any analysis that does justice to metaphor in general and to specific examples of metaphor must have the resources to distinguish among these variables.

In addition, once we have distinguished these seven ways in which metaphors can differ among themselves, we can also see that metaphors may overlap with literal utterances in many of these ways as well. So, for instance, we saw in chapter 4 that stories intended literally, whether as fact or as fiction, can also be intended to generate aspects in my sense of the term, and that those aspects can carry significant conversational weight, can be novel and resonant, and can be something to which the speaker is fully committed. One major difference between metaphorical and literal utterances which generate aspects is that, while the aspects generated by all such literal utterances are internal, the aspects on which many — though by no means all[37] — metaphors rely are external. Another major difference, related to the previous one, is that with metaphor speakers are much less commonly committed to the overall subject characterization imposed by the aspect than they are for literal utterances. In general, the only essential and universal difference between metaphorical utterances and those literal utterances that generate aspects is that in speaking metaphorically, the speaker does not primarily intend to commit herself to the semantic content she utters, while with a literal utterance she does so intend. Our discussion of putative 'thrice-apt' metaphors raised the possibility that for some special cases the speaker *does* intend to commit herself to the sentence's semantic content after all, but also still intends her utterance metaphorically. I suggested, though, that these cases may be better treated as instances of symbolism than of metaphor.

So I think that distinguishing these seven variables among metaphors allows us to get a more precise understanding both of the overall topography of metaphor itself, and of the relations between metaphorical and literal speech. In §1.4 I complained that theorists of metaphor tend to focus exclusively

---

[37] Recall from §4.5.1 that metaphors like 'Bill is the quarterback of our office', or 'Alice is the mother of our friends' rely on internal aspects.

on just one end of the continuum from 'ordinary' to 'poetic' metaphor. Pragmatists tend to focus on

ordinary conversational metaphors, and dismiss the rich resonant effects of more poetic metaphors as

mere by-products. Theorists who invoke aspects and 'interactionism', on the other hand, typically focus

on rich, poetic metaphors, precisely because these cannot be fully treated in propositional terms. It's easy

to see what the narrow pragmatists miss; much of my discussion of aspects has aimed at addressing just

this neglect without thereby abandoning a pragmatic theory altogether. However, focusing on aspects as I

have then creates a corollary difficulty for me, in explaining why ordinary conversational metaphors like

(27) and (33) *are* still metaphors after all. I turn to this difficulty now.

The problem for an aspect-based account of ordinary conversational metaphors is that the

relevant aspect doesn't appear to do any work in these cases. We don't need to explore just how F

interacts with a's characterization to discover which a-features can be tightly matched to F-features,

because the relevant F-features — regardless of what particular subject the aspect is being applied *to* —

have already been established by repeated use. Stern calls these cases '**routinized**' metaphors, which for

him means that their particular context of application is no longer crucial in determining the content

meant.[38] William Gass describes the process by which metaphors become 'routinized' thus:

> Each time we refer to someone as a mouse we beat the same path through the concept until at last this path
> is broad and movement is easy and immediate. The sense of traveling through strange lands is lost, and
> when we no longer have to hunt for the point of the comparison, we begin, quite justly, to wonder why
> Clifford wasn't called shy and frightened in the first place; for, of course, Clifford could have been a
> mouse for every reason, his whole life seen through that system; but this particular metaphor has slid even
> beyond proposition to comparison — where the metaphor says, in effect, Clifford is like a mouse because
> both are afraid — until it is nearly a case of catachresis, which has little or no figurative commitment, the
> word 'mouse' now being wrongly used to mean someone shy and easily frightened.[39]

I think Gass is correct that metaphors can die through repeated use and thereby create new literal

meanings — although, as we saw in §5.4.3, familiarity does not always cause this slide toward

literalization. The crucial claim, though, is that for 'routinized' metaphors like (27), (33), and now

(40) Clifford is a mouse,

[38] Stern, *Metaphor in Context*, p. 311.
[39] "In Terms of the Toenail: Fiction and the Figures of Life," p. 67.

the path from what is said to what is meant becomes so 'well-trodden' that aspects do no work at all in them. Is this true?

I don't think it is true for examples of this sort. It is more true about some other examples, which I will turn to next, but for those cases we have much weaker intuitions that the utterances are indeed metaphorical. Thus, we have good reason to think that whether aspects still play a role in comprehension is indeed criterial for an utterance's being metaphorical. Of course, it is true that in these routinized cases we don't have to go through any conscious process of teasing out the aspect's application and the consequent *M*-features, as we do have to do with a wild, novel metaphorical utterance like

(42) She was false as water.[40]

But as we saw in §1.1.2.1, in my discussion of Recanati's 'p-literality', conscious awareness is an inaccurate gauge in many other cases of pragmatic processing and in the study of mental phenomena generally, and so we shouldn't rely upon it here.

I think that the work aspects do still do in these 'routinized' cases is manifested in at least two ways. First, although the content to which the speaker commits herself by her utterance may be quite restricted and routine, she usually still intends to *suggest* at least some of the features and especially the attitudes that would be delivered by actively applying the aspect. I think this is true even of the relatively disposable cases above. To the extent that this is true, the aspect's effects outstrip the simple determination of *M*-features, and the literal substitute is not equivalent after all.

Second, equally routine conversational utterances which exploit the same term metaphorically can still express very different contents, depending on the speaker's intentions and the subject under discussion. Thus, 'is a pig' can be used metaphorically to characterize someone as greedy, gluttonous, ugly, sloppy, lustful, or uncivilized, or some combination of these; and 'is a mouse' can be used metaphorically to characterize someone as timid, shy, nervous, quiet, or ineffectual, or some combination. But the only salient quality that these potentially predicated features have in common is being prominent elements of our characterization of pigs or of mice, respectively. Which of the features is being applied to

---

[40] *Othello*, V.ii.134.

the subject will depend on the particular subject characterization and on which of its features are most prominent in the conversational context. Thus, the determination of the metaphor's content *does* seem to depend upon finding matches between the governing characterization and the particular subject characterization after all. The aspectual means that I have described for moving from what is said to what is meant, which is more obviously required for the richer, more 'poetic' cases, provides the most economical explanation in these cases as well.

One might object to this that the fact that multiple distinct contents can be delivered by equally routine metaphorical utterances of the same term only reveals that there's a greater range of semantic ambiguity or polysemy in the meanings of the word 'pig' or 'mouse' than we might have expected: instead of *one* dying or dead metaphorical meaning, we have several. On this view, the relevance of the subject and the speaker's intentions is the same as for the resolution of any ambiguity or polysemy, but this doesn't show that anything distinctively aspectual — or even pragmatic — is going on. In response to this, note that the putative additional meanings of 'mouse' or 'pig' are related to one another in two ways which make either an ambiguity or polysemy thesis difficult, though not impossible, to maintain.

First, the various meanings can all apply to the same basic sort of thing: to human beings (and even to human beings who are much more specifically of the same sort: middle-aged white-collar married men, for instance). This distinguishes the additional meanings for 'mouse' or 'pig' from typical cases of ambiguity, like 'cape', 'mass' or 'bank'. In these latter cases, it's criterial of each of the different meanings that they apply to a certain category of thing, and the different meanings apply to very different categories.

Second, the various putative meanings for 'pig' or 'mouse' predicate properties which are truth-conditionally very different. This distinguishes the additional meanings from typical cases of polysemy, like 'fast' or 'healthy'. Although what's involved in being fast is manifested differently for drivers, cars, roads, tires, and food, these are all essentially related, albeit in different ways, to a single defining feature:

of *moving quickly*.[41] Fast drivers count as fast in virtue of moving quickly; fast cars count as fast in virtue

of enabling drivers to move quickly, and so on. For routinized metaphors, by contrast, the different ways

of being metaphorically mousy or piggish are much more various: they are not just different

manifestations of a single defining feature. The only thing that unites them is their relation to the relevant

governing characterization, but this relation is not definitional. That is, Bill is (metaphorically) a mouse

in virtue of being timid; Clifford is (metaphorically) a mouse in virtue of being nervous; David is

(metaphorically) a mouse in virtue of being ineffectual, and so on. These properties are quite different,

and none of them are definitionally *ways* of being a mouse, as moving quickly and enabling to move

quickly are definitionally ways of being fast.

All that being said, it is of course true that the range of features available for possible application

by a routinized metaphorical use of 'is a pig' or 'is a mouse' is much narrower than the potential range for

a rich, novel metaphor. Routinized metaphors *are* routinized, and so the number of relevant elements in

the governing characterization is considerably reduced. In these cases, the characterization may be no

more than a crude stereotype. The aspect thus does do less work in these cases: there is less re-

organization and less selection among many candidate elements. Gass does a good job in the passage

quoted above of evoking the process by which the aspect becomes limited in this way. (Although, as we

saw in §5.3.3, familiar aspects can still be reinvigorated into novel metaphors as well.) But routinization

is not yet literalization, nor does it involve the complete extrusion of the aspect from the interpretive

process.

When a 'routinized' metaphor actually does die into the new life of literal meaning, no such

variability remains; the aspect ceases to play any determining role. But in these cases, our intuitions also

---

[41] The range of variability for even routinized metaphors is *a fortiori* broader, and of a different kind, than the sort of variability that Charles Travis claims is manifested in, say, the different ways that something can be green: putting aside vagueness about the boundaries of greenness, Travis points out that a green apple is green only on the outside, while green jade must be green through and through, and that a green apple is naturally green, while a green car has green paint applied to it. Here just a single feature, of being a certain color, is relevant. See his "On Constraints of Generality"; see also Searle, "Consciousness, Intentionality, and the Background." Travis, but not Searle, takes this variability to undermine the tenability of the very idea of literal meaning.

support the claim that the utterance is no longer metaphorical. Examples of such genuinely dead metaphors are

> (43) The table has four *legs*.
> (44) Dartmouth is located at the *mouth* of the river Dart.

With truly dead metaphors, although we think we discern some etymological trace of metaphoricity, we have no qualms about directly disquoting from an utterance to an attributed belief, or about directly assessing the utterance's truth-conditions. There are also many cases that fall somewhere in between being fully dead and being heavily routinized, such as

> (45) John is a rising star of the New York political scene.

For these cases, we have very weak or mixed intuitions about metaphoricity and some hesitation, perhaps rather slight, about directly disquoting. And in such cases, we find very little variability in the term's application, though there may still be some minor variability. Lakoff and his colleagues specialize in cases of this sort:

> (46) It's difficult to *put* my ideas *into* words.
> (47) The idea is *buried in* terribly dense paragraphs.[42]

are said to reveal the underlying metaphorical schema LANGUAGE IS A CONDUIT; while

> (48) We've got the *framework* for a *solid* argument.
> (49) He is trying to *buttress* his argument with a lot of irrelevant facts, but it is still so *shaky* that it will easily *fall apart* under criticism.[43]

depend upon the schema AN ARGUMENT IS A BUILDING. The fact that similarly routine applications of the same purportedly metaphorical terms do not express very different contents suggests that these uses of the terms are metaphorically moribund if not entirely dead. I think it's important to admit, however, that the boundary between the metaphorical and the literal is fuzzy, and shifts with time and use. This should come as no great surprise, because literality and semantics are a matter of convention, and conventionality in the realm of language is itself a somewhat vague matter.

---

[42] *Metaphors We Live By*, p. 11.
[43] *Metaphors We Live By*, p. 98.

5.4: *Contrasting Accounts of Aspectual Thought and Metaphorical Comprehension*

At long last, my own account of metaphor is complete. I have specified the pragmatic means by which speakers manage to communicate metaphorically with their hearers: I claim that they intend for their hearers to cultivate an aspect in thought, in order to determine the primary content to which the speakers intend to commit themselves. That aspect is generated in a more or less direct fashion from the sentence uttered (§5.2), and it then puts the hearer in a position to identify the features of the subject which are most tightly matched to the most prominent features of the governing characterization (the *M*-features). The aspect's cultivation produces metaphor's vaunted and maligned 'non-propositional' effects, while the identification of the *M*-features provides the speaker's intended metaphorical meaning. I have also specified seven of the most important ways in which metaphorical utterances differ from one another and from literal speech, both in terms of the aspect cultivated and in the content determined (§5.3).

In this section I compare my view explicitly with those of several other philosophers we have encountered earlier. I hope that this lends additional credence to my view, by showing how my account builds on those existing theories. More importantly, this will show that my theory can reconcile within a single, unified account the various phenomena we encountered in chapter 1 — phenomena which other theorists have assumed to require disparate theories of metaphor.

5.4.1: *Black's Interactionism*

My view of metaphorical comprehension is both a generalization and a refinement of Max Black's *interactionism*: the idea that metaphorical meaning depends upon some sort of 'interaction' between the 'topic' and 'vehicle'. It is a generalization because I claim that characterizations and aspectual thought are relevant, not just in the comprehension of metaphor, but in other contexts as well: at least, in psychological explanation, in comprehending historical and fictional narratives, and in aesthetic appreciation. In chapter 3, I presented my view of characterizations as a refinement of Black's discussion of metaphor as involving the interaction of two 'systems of associated commonplaces.' I specified the

way in which these 'commonplaces' are structured into a system (by relative assignments of prominence, centrality, and depth); the requirement on a thought being included within or 'associated' with the system (fittingness); and the sorts of thoughts that can be included (not just 'commonplaces', but other features as well, such as highly specific experientially represented properties and affective attitudes).

Besides the differences which result from these differences between our respective 'systems', the most basic difference between my view and Black's is that I locate the relevant sort of 'interaction' on the level of thought rather than the level of language, as Black (usually) seems to do. As a result, I can make more systematic sense of the way in which the 'systems' associated with very different parts of speech — names, kind terms, descriptions, and predicates — can 'interact' in the same basic way to produce aspectual thought. At the same time, I am free to acknowledge the important associations between words and characterizations which led Black (and with him, Eva Kittay, Andrew Goatly, and Patti Nogales) to a semantic view.

If we turn to aspectual thought specifically, my view is a form of 'interactionism' because, as we saw in §4.5.2, which aspect is produced depends neither upon just the governing characterization, nor on just the subject characterization, but on the interaction between them. More specifically, on my view, as on Black's, "the metaphor works by applying to the principal subject a system of 'associated implications' characteristic of the subsidiary subject."[44] I go into greater detail than Black does about just how this 'application' works. In particular, the process of matching which I specify makes good on the metaphor of 'interaction' in a much more concrete and precise way. On my view, it's not just that appropriate features of the subsidiary subject (that is, governing characterization) can themselves be applied to the principal subject (that is, subject characterization), or that when the same features occur in both subjects those features are then pushed to the fore and unmatched features are 'filtered' out, as Black maintains. In addition, the matching process can alter the structural role of features of the subject which are merely similar in some respect to features of the governing characterization.

---

[44] "Metaphor," p. 44.

5.4.2: Searle's Principles

My account can also be seen as a version of Searle's pragmatic theory of metaphor. Recall from §1.2 that Searle specified as his task the same aim that I articulated: to identify a specific "shared system of principles" on the basis of which we can "provide a rational reconstruction of the inference patterns that underlie our ability to understand" metaphors.[45] He also accepted that we should state these principles "in a way that makes it clear how metaphorical utterances differ from other sorts of utterances in which speaker meaning does not coincide with literal meaning."[46]

His solution to this challenge is to offer eight principles by which the range of possible metaphorical meanings is determined on the basis of the metaphorical word uttered. Those principles appeal to, among other things, 'salient defining characteristics', 'well-known properties', 'facts about our sensibility', and similarity in types of the conditions specified. The broad range of possible meanings identified by these principles is then narrowed down to the speaker's intended meaning, by appealing to only those candidate meanings which specify possible features for the subject itself.

We might think of my account as replacing Searle's eight principles for generating possible metaphorical meanings with what I call the governing characterization associated with the metaphorical term. My governing characterizations do also appeal to 'salient characteristics,' 'well-known properties,' and 'facts about our sensibility,' among other things. These are the sorts of relations which can lead a feature to be included in a characterization. But I have tried to motivate appeals to such relations by grounding them in our having a single, unified sort of mental representation which does independent explanatory work in a theory of the mind. Further, I have tried to build into those mental representations the resources for explaining, not just how the propositional contents of the speaker's meaning can be derived from the semantic contents of her utterance, but also how metaphors can bring about the non-propositional effects of aspectual thought which so many people take to be central to the phenomenon of

---

[45] "Metaphor," p. 104.
[46] "Metaphor," p. 78.

metaphor. ==By contrast, Searle treats all non-propositional effects as inherently beyond the purview of a theory of meaning.==

Searle's principle for narrowing the broad range of potential meanings down to the speaker's intended meaning finds its analogue in my view in the role that the subject characterization plays in guiding metaphorical interpretation. As I argued in §4.5.2, when we cultivate an aspect in thought, we do not just notice the most prominent features of the governing characterization itself. Rather, we attend to those features in the subject characterization which can be matched to prominent and central features in the governing characterization. The process of matching as I've articulated it does rely on finding similarities between the two sets of features, and in this respect my account may seem to repeat the vague invocations of similarity or resemblance by Grice, Bach and Harnish, and to some degree Searle himself. But I have gone beyond the vague invocation by specifying four specific sorts of similarity that can be exploited.

Finally, even though I have tried to show that metaphor relies on a sort of mental representation which does independent explanatory work in a theory of mind, I think I have also identified what is distinctive about metaphorical communication as opposed to other non-literal and indirect utterances, as Searle rightly insists an adequate theory of metaphor must. In speaking metaphorically, I claim, a speaker intends for her hearer to determine the content of her primary communicative commitment by using an aspect to identify the subject's putative *M*-features. (One might object that I have not isolated a distinctive feature of metaphor after all, because similes rely on much the same cognitive mechanism. ==However, for this very reason I share the view those who think that simile is a sort of metaphor:[47] I take it that the 'like' or 'as' of a simile functions as a sort of 'hedge,'[48] weakening the strength of the speaker's commitment to the governing characterization's applicability as an aspect for thinking about the subject.)== Metaphors differ in principle from literal utterances in which the speaker also intends for the hearer to

---

[47] Cf. Ortony, "Beyond Literal Similarity," Fogelin, *Figuratively Speaking*. Both Ortony and Fogelin thinks that metaphors are elliptical similes, with the 'like' or 'as' present at the level of logical form. However, they both also think that similes and metaphors both rely on a distinctive sort of *figurative* similarity, and so they think that the two tropes exploit the same basic mechanism.

[48] Cf. George Lakoff, "Hedges: a Study in Meaning Criteria and the Logic of Fuzzy Concepts."

cultivate an aspect in the following way: only in metaphor does the cultivation of the aspects serve as the means for identifying the content of the speaker's primary communicative commitment. As I said in §5.3.5, there also tend to be other differences between the aspects employed in the service of metaphorical communication and those employed in connection with literal utterances — for one, the aspects involved in metaphor tend to be external — but these differences are not differences in principle.

5.4.3: *A Davidsonian 'Noticing' Account: of Aspects, of Metaphorical Content*

Davidson's view of metaphor has served as both an inspiration and a foil throughout the development of my account. We saw in §1.3 that Davidson seems to be driven to his non-cognitivism about metaphor precisely because he assumes a clear and exhaustive division between rational, inferential laws of thought and merely causal, associative processes. In addition, this assumption also leads him to his positive view of aspectual thought and of metaphorical interpretation. On his view, when we think of one thing as another, we are caused to notice surprising similarities between the two things. This in turn causes us to attend to otherwise neglected features of both things, but in particular of the metaphor's subject. Although Davidson doesn't explicitly say this, presumably he can also maintain that noticing those similarities also causes us to transfer some of our affective attitudes about the one thing onto the other.

Davidson's view is thus a fairly pure version of what I have called a 'thought *plus* thought' or juxtaposition account of aspectual thought. It's just the sort of account you would construct if you wanted to explain aspectual thought using minimal, and merely causal, psychological resources. But we have seen at multiple points that neither a merely causal account nor a juxtapository account (whether merely causal or not) can account for how aspectual thought actually works.

First, as it stands, Davidson's story has no resources for explaining the power of aspectual thought for genuine discovery: its ability to make us aware of previously unknown features, as opposed to known but neglected and nearly-forgotten ones. In perceptual 'seeing-as', we are confronted with some actual thing which we can examine closely. So in that context, we may possibly discover new features of

it. But in aspectual thought, the subject is not typically present to be examined; all we have to rely upon are our own thoughts about it. And here there can be no undiscovered features, unless these are simply unnoticed consequences of beliefs we did already possess. But aspectual thought, especially in the context of metaphor, does often lead us to think that the subject possesses new features.

Davidson might try to explain this by taking something like my own view on board here. By virtue of noticing certain similarities between the two things, he might say, we are caused to *wonder* whether the one thing also possesses certain of the other thing's features. We can then go out and investigate whether the subject does have that feature after all. To the extent that Davidson does appeal to something of this sort, the distinction between his merely causal account and mine will begin to narrow, because he will be appealing to a system of thought associated with one thing as a cognitive *tool* for thinking about the other. Further, within the context of metaphorical communication, we treat the fact that the speaker has made a speech act which commits her to the claim that the subject possesses some new feature to be a *reason* for thinking that the thing in fact does have that feature. And more than a merely causal story is required to account for this.

But even just taking this much of my view on board is still not sufficient to explain how feature introduction works. I have argued that there are more specific and complex relationships between aspect application and feature introduction. We don't wonder whether the one thing possesses just *any* of the other thing's features. Rather, we only concern ourselves with features which would be *fitting* for *a*, given the rest of what we know about its character. And we aren't confined to wondering whether *a* possesses just the same features as *F*-things have. We wonder about *a*-features which would be matched in certain sorts of systematic ways (by scalar, analogical, or metaphorical similarity) to known *F*-features. Finally, we often take the fact that a certain analogy obtains between some of *a*'s and *F*'s features itself to be a reason for thinking that *a* has a certain further feature. We reason in this way even in serious scientific contexts. Once again, the fact that we rely on intuitions about reasons in this context means that a merely causal explanation is insufficient.

In addition to being unable to explain feature introduction, Davidson also can't explain the more general and ongoing utility of aspectual thought in contemplation.[49] As we saw in §4.6, aspects are effective tools not just for getting others to notice surprising features, but also for organizing and negotiating our way among our own thoughts and feelings about something. Davidson's account simply can't make sense of this. As Black says, "What then, on Davidson's view is a soliloquizing thinker, using metaphorical language, supposed to be doing? Nudging and provoking himself to pay attention to some covert likeness? But surely he has already done so?"[50] Otherwise, Black seems to be suggesting, why would the soliloquizing thinker think to invoke that aspect in the first place? Davidson might respond that a thinker soliloquizes metaphorically (or thinks aspectually) because he has a *suspicion* that this may nudge him into noticing neglected features. I think this often does happen. The point remains that the aspect may continue to be cognitively useful for us as an organizational tool even after all the nudging is over.

Davidson does recognize that metaphors have *some* continued effect even after our first encounter with them. In his criticism of "second meaning" theories of metaphor, he writes: "Novelty is not the issue...What we call the element of novelty or surprise in a metaphor is a built-in aesthetic feature we can experience again and again, like the surprise in Haydn's Symphony No. 94, or a familiar deceptive cadence."[51] But this sense of 'built-in surprise' as an *aesthetic* feature does not capture the sense of continued cognitive utility and interest which Black and I are pointing to here. Besides, merely saying that there *is* a 'built-in surprise' does nothing to explain why this should be so. On our first encounter with a metaphor, the surprise is supposed to be caused by the noticing of new similarities between the two things. But once we know those similarities are there, why should they surprise us?

So far, I have been discussing Davidson's theory of aspects, because this is all he thinks there is to a theory of metaphor. But we can see even more clearly what is wrong with a Davidson-style, causal

---

[49] This deficiency parallels (and perhaps is the product of) a deficiency in Wittgenstein's own account of seeing-as, which he restricts to the dawning of an aspect.

[50] "Afterthoughts on Metaphor: How Metaphors Work: A Reply to Donald Davidson," p. 188.

[51] "What Metaphors Mean," pp. 252-3.

juxtaposition account of aspectual thought if we bracket Davidson's non-cognitivist commitment for a moment. If a philosopher were willing to grant that metaphorical utterances have propositional content, but also wanted to take seriously the importance of 'seeing things in a new light' for metaphor, the initially most natural way to go would be to exploit a Davidsonian view of aspects as the means for determining content. Specifically, as I mentioned briefly in §1.3, one could claim that speakers who speak metaphorically intend to predicate of their subjects just those features that the aspect 'nudges us into noting.' The fact that Davidson's view of aspects won't work even as one component within a theory of metaphorical content, rather than as the whole of a theory of metaphor, as he intends it, is an additional mark against it.

I take it that in my terms, the features of a which we are 'nudged into noting' are just the features in the a-characterization that are most increased in prominence by applying the aspect. And often, my M-features — that is, the a-features which are most tightly-matched to the most prominent and central F-features — will also be those a-features which are most increased in prominence by applying the aspect. After all, the a-features which find matches are just those which get increased in prominence, on my view. So in many cases, the simple Davidson-style account of aspects will deliver intuitively appropriate contents for metaphorical utterances. However, being tightly-matched to an F-feature and being increased in prominence need not go together. When they come apart, only the tightly-matched features enter into the speaker's meaning. Thus, my account delivers the right answer when the two theories conflict, as they sometimes will.

Being tightly-matched and being increased in prominence can come apart in two ways. On the one hand, some feature Q in a's characterization that is not matched to any F-feature may end up being rendered prominent because of other changes wrought by the aspect-application. Even if the hearer notices Q's applicability to a, and even if the speaker *intends* for him to notice this, Q should intuitively not be included in the speaker's metaphorical meaning — although it well may be part of a *further* indirect speech act of suggesting or insinuating. But on the 'noticing' account of aspect application, Q would count as part of the content. So, for example, Alice might say to Bill

(50) Charles is a strutting rooster.

She might intend as one effect of her utterance to get Bill to notice Charles's tendency to recount stories repeatedly in which he figures as the hero, even when others know that he was really just a minor player. She might also thereby intend to get Bill to notice that Charles's latest tale is probably false. On the simple 'noticing'-based story, this latter claim should be part of the content of Alice's metaphorical assertion. But intuitively, Alice has only claimed that Charles is a certain roosterish sort of preening braggart. She has not claimed, but merely insinuated, the further thought that this story of his is false. On my theory, because there's no element in the characterization of *being a strutting rooster* that matches the tendency to lie, the claim that Charles' latest story is probably false is no part of what Alice asserted.

On the other hand, some feature R in a's characterization that *is* tightly matched to a feature in F's characterization may already be mutually known to be highly prominent for independent reasons. If this is so, then R will not be measurably increased in prominence by the aspect's application, and it certainly will not be noticed for the first time as a result. But intuitively, the predication of R to a should still be part of the metaphorical content. So, for example, two men might be discussing the strikingly red hair of a woman sitting across the café, and they might begin competing to find the best characterization of it. When one of them says

(51) Her hair is bands of scoured copper,

he is claiming that her hair is *that* shade of red and has that surface texture. This remains true even though the utterance didn't lead either of them to notice anything new about her or her hair.

5.4.4: *The 'Complicity Effect'*

In the next two sections, I tie up some important loose ends from chapter 1 by showing that my account can explain two features of metaphor which have been taken to demonstrate the inadequacy of a pragmatic account. Indeed, my pragmatic account explains these features in a more parsimonious way than the theorists who first raised these features can.

Recall from §1.3 that some philosophers — notably Wayne Booth, Ted Cohen, and Richard Moran — have postulated at least a non-cognitive component or dimension for metaphor on the grounds that metaphors have a mysterious compulsory rhetorical force. That is, for at least some metaphors, even a hearer who strongly disagrees with what the speaker means is nonetheless 'implicated' in the metaphor simply by his comprehension of it. As Cohen says, a metaphorical insult or a cruel joke "is all the more painful because the victim has been made a complicitor in his own demise."[52] Moran claims that this is quite different from "an ordinary literal insult," for which "'getting the point' does not interfere with denying or repulsing it."[53] Wayne Booth formulates the point about metaphor's 'complicity' and 'irresistibility' thus:

> The speaker has performed a task by yoking what the hearer had not yoked before, and the hearer simply cannot resist joining him; they thus perform an identical dance step, and the metaphor accomplishes at least part of its work even if the hearer then draws back and says, 'I shouldn't have allowed that!'[54]

Indeed, we might add that the effect seems to survive the *speaker*'s own disavowal of the metaphorical comparison. Thus, an utterance of

> (52) I would not describe my opponent as a scavenging hyena, for that would not be true. But I will say that questions have been raised about his ability to come up with his own ideas.

carries much, if not all, of the same force as an outright assertion of the initial metaphor. But if, as these theorists maintain, it is in the very nature of cognitive, propositional communication that mere understanding cannot compel assent, then metaphor must have some essentially non-cognitive, non-propositional dimension.

I can now explain why metaphors have this 'complicity effect'. It is a consequence of the role that characterizations and aspects play in the process of metaphorical comprehension. Specifically, the hearer's very ability to discern the speaker's meaning shows that he can cultivate the relevant aspect. That is, it is only by way of applying F as an aspect to a that the hearer can determine which features the speaker intends to predicate of a. If the hearer couldn't apply the aspect, he would respond with sheer incomprehension rather than disagreement. But the hearer's ability to cultivate the aspect in turn reveals

---

[52] "Metaphor and the Cultivation of Intimacy," p. 10.
[53] "Seeing and Believing: Metaphor, Image, Force," p. 91.
[54] "Metaphor as Rhetoric: The Problem of Evaluation," p. 52.

that he has uncovered appropriate matches between features in the two characterizations. And *this* in turn reveals that he has at least temporarily adopted subject and governing characterizations sufficiently like the speaker's own. But one or the other of these characterizations is objectionable to the hearer, and so he is complicit by virtue of adopting a characterization he objects to.

Because of all this, to reject an assertive metaphor entirely, it is not sufficient just to deny that *a* is *F*, because the speaker obviously didn't intend to claim *that*. Nor is it sufficient to deny that *a* possesses the intended *M*-features, because this will still leave the aspect itself unchallenged. In this sense, ==**disagreement** about the content is still disagreement *within* the context of the aspect's application.== To reject the metaphor entirely, one must either respond with incomprehension, or else concede that one knows what the speaker is 'getting at', but insist both that *what* she is getting at is false, and further that the *way* in which she is getting at it relies on an inaccurate and objectionable characterization. (One effective way to do this is to follow up with a metaphor which relies upon a very different, and presumably more appropriate, aspect.)

But with these points on the table, the 'complicity effect' should seem both less mysterious and less distinctive of metaphor. It is a special, and more complex, instance of the much more pervasive phenomenon of conversational presupposition. An assertion's presuppositions are, by definition, those assumptions that both the assertion and its denial share, and so that persist unchallenged through a denial of the assertion's content. In order to comprehend the speaker at all, the hearer must determine which presuppositions she is assuming. This already requires some cooperation with the speaker, although it does not require the specific, nuanced understanding involved in cultivating an aspect. Because presuppositions persist through denial, it can be difficult to find an appropriate response to even a direct, literal speech act whose presuppositions one rejects. (Two typical examples of the latter are the presuppositions of existence and uniqueness triggered by a definite description and the presupposition of a past continuous activity of *F*-ing triggered by 'X stopped *F*-ing'.) But this presuppository power is neither magical nor limitless. A hearer can and should 'explicitate' and then explicitly reject

presuppositions which he finds objectionable. Only then can the conversation proceed as a cooperative activity involving both participants.[55]

It is true that metaphors' presuppositions are much more difficult to 'explicitate' and so to disavow than ordinary conversational presuppositions. This additional difficulty has at least four sources, in addition to the crucial point about the role of aspect-application in deriving content above. First, ordinary literal speech acts often presuppose just a single or a few propositions, while metaphors rely on characterizations, which consist of many elements related in a complex structure. Second, we lack an established vocabulary for making our characterizations explicit. Specifically, without a well-developed vocabulary for fittingness, the hearer has no convenient means for challenging this sort of applicability, which can remain in place even after truth is disavowed. Third, characterizations are interpersonally less stable than ordinary conversational presuppositions. This makes it less determinate just what the speaker must be assuming, and so leaves the speaker with more wiggle room ('I didn't mean to assume *that*! I only meant...'). Finally, affective attitudes play a more complex and essential role in characterizations than in many ordinary conversational presuppositions. In particular, because affective attitudes are permeant, they can continue to color other features in a characterization even after the features that initially underwrote them have been rejected.

Even these differences, though, are not essentially distinctive of metaphor. Characterizations and affective attitudes are closely, even conventionally associated with many terms in the language. In this sense, much of our use of language has a non-cognitive dimension. For this reason, Moran is wrong to draw such a sharp distinction between the deniability of metaphorical and 'ordinary literal insults'. In particular, whenever a speaker employs a derogatory (or laudatory) term associated with a characterization that the hearer rejects, such as 'spic' or 'fag', the hearer will be in the same uncomfortable position of 'complicity': merely denying the applicability of the term to the subject under

---

[55] Cf. Brandom, *Making it Explicit*. My thinking here has also been influenced by Josef Stern's claim that presupposition plays an important role in metaphor; see *Metaphor in Context*, especially chapters 4 and 5.

discussion leaves the objectionable characterization in place as part of the common conversational background.

There is one important way in which the essential role that aspects play in generating the 'complicity effect' for metaphorical utterances distinguishes them from other utterances that also generate 'complicity'. The effect is still triggered with a metaphorical utterance when the hearer accepts the governing characterization, and admits that it *could* have some appropriate application, but just wants to deny its applicability to the particular subject under discussion. In a case like this, unlike with an objectionable derogatory term, the hearer is willing to accept that *someone* might aptly be described as, say, 'a damp kitchen sponge'; he just wants to deny that *this* subject is aptly describable thus. Here, the fact that the hearer can apply the aspect at all makes his denial ring hollow in a way that doesn't happen if, say, he denies that the subject is 'a lazy bum' (supposing he agrees that some people *are* lazy bums). However, even in a case like this, we don't need to postulate any distinct and essentially non-cognitive dimension, *pace* Moran. Rather, the 'complicity effect' here depends on the facts about the *process* of metaphorical comprehension described above.

The claim that the aspect does play a special role in generating the complicity effect in metaphor is supported by the fact that the effect becomes weaker as the metaphor becomes more highly conventionalized. So, merely denying the claim made by an utterance of

   (53) My opponent is a mouse

leaves the subject less tainted by the aspect, and the hearer less embarrassed by his understanding, than does a mere denial of the claims made by

   (54) My opponent is a damp kitchen sponge

or

   (55) My opponent is a tailwagging lapdog of privilege.[56]

This is because a more conventionalized metaphor requires less subject-specific work in cultivating the aspect's application: the hearer is already familiar with the basic outlines of its application in advance. As

---

[56] This insult is also from Moran. "Seeing and Believing: Metaphor, Image, Force", p. 90.

a result, his comprehension of the speaker's meaning doesn't necessarily reveal that he has grasped the speaker's specific intended matches between the governing and subject characterization.

5.4.5: *'Lodging in the Words'*

Moran, Booth, and others took the 'complicity' effect to demonstrate that metaphor had an essentially non-propositional component; my theory explains this phenomenon as a manifestation of the more general phenomenon of conversational presupposition. It turns out that much the same phenomenon of conversational presupposition also explains the features that have been taken to establish the very different claim that metaphor is a semantic phenomenon. Recall from §1.1.2.2 that Hills, Bezuidenhout and Stern pointed to two important ways in which metaphor differs from typical cases of **implicature**. **First**, in speaking metaphorically a speaker makes a speech act whose content is available for agreement or disagreement, by saying something like 'Yes, it is', or 'No, it's not'. **Second**, later uses of the same words will most naturally be taken to express the same, metaphorical content. That is, metaphorical speech sets up a 'default' interpretation within the conversational context. Hills and Stern claim that this demonstrates that metaphors alter the contents of the *words* uttered in that context. I argued in §1.1.2.2 that the interpretive options are not as limited as Hills and Stern make them out to be, and so that a semantic explanation is not warranted. However, I still need to offer my own explanation of these phenomena.

Hills's first point was that other speakers can respond directly to a speaker's intended metaphorical content. Thus, Benvolio can easily respond to Romeo's utterance of

(56) Juliet is the sun.

by saying

(57) Yeah, I know. Juliet is quite beautiful and wonderful.

But, Hills claims, this is not true of the paradigmatic Gricean implicature,

(58) Mr. X's command of English is excellent, and his attendance at tutorials has been regular.

In this case, it's not normally appropriate to respond by saying

(59) Yeah, I know. Mr. X is really not a viable candidate.[57]

Notice, though, that this contrast distinguishes (58) not just from metaphor, but from all figurative utterances. So, for instance, Bill might agree with Alice's sarcastic utterance of

(60) Jane's really been a fine friend to me in these last few weeks

by saying something like:

(61) Yeah, she really let you down.

Or he might disagree with her utterance by using her same words sincerely:

(62) She *has* been a fine friend!

But we are much less likely to treat sarcasm as a semantic phenomenon, because words would then mean both what they do *and* their 'opposites'.

I think the explanation of this general phenomenon depends on what the speaker is obviously interested in committing herself to. In many cases of implicature, including Grice's letter of recommendation case, the speaker first makes a speech act in which she undertakes some sort of commitment to the content of the sentence construed literally, and then makes another, indirect speech act as well. By contrast, when a speaker uses her words figuratively, she has obviously made a speech act which commits her to *some* content, but it is also obvious that she isn't interested in committing herself to the contents literally expressed by her words. The degree to which the initial speech act is directly relevant to the conversation determines the degree to which it is a natural subject for response.

The 'paradigmatic' Gricean implicature is thus actually rather atypical in this respect. For many other cases of indirect speech or implicature, it *is* felicitous for the hearer to take up and respond to the speaker's indirectly expressed content.[58] So, after the following exchange

(63) Alice: Do you want to go see a movie tonight?
    Bill: I have an exam tomorrow,

it is perfectly felicitous for Alice to respond to Bill's *implicated* claim: that he cannot go because he has to study. For instance, she might say

---

[57] I share Hills's intuition about the two cases here, but several readers have suggested that (59) might be an appropriate response after all. If so, so much the better for my case.
[58] Cf. Searle, "Indirect Speech Acts."

(64) But you promised we'd hang out this week, and this is the only possible night.

In (63), Bill's actual response does not answer the question if it is just construed on its own: Alice hasn't asked what he is doing tomorrow. But his implicated response does answer the question, and his literal response implicates what it does because it serves as a reason for the implicated answer. Bill's implicated answer is the obvious linchpin in the exchange, the indirect step which makes what he says relevant. In the Gricean case of (58), by contrast, the speaker has taken care to fulfill the letter of the request in his actual response. He *has* directly answered the question actually asked: 'What can you say to recommend Mr. X?'. The implicature that Mr. X is not a good philosopher is generated precisely by the speaker's cagey refusal to make any sort of explicit commitment on the matter that's clearly but not fully explicitly at hand, which is Mr. X's quality specifically *qua* philosopher.

Although for metaphorical utterances (and figurative utterances generally) it is quite natural for a hearer to respond to the speaker's intended, figuratively interpreted content, it is also important that the hearer is not compelled to respond in this way. As I argued in §1.1.2.2, the speaker has still *said* the sentence with its literal content, and that content is available for response as well. Although the speaker didn't want to be taken as committing herself to that content, her hearers have the right to take it up and challenge it. Because it is obvious that she didn't mean what she said, a charitable, fully cooperative hearer will not respond in this way. But the speaker must rely upon the hearer's indulgence here as she wouldn't need to if she had meant what she said.

This much explains the naturalness of responding to the speaker's intended metaphorical *content* rather than the literal content of what she has said. But Hills and especially Stern are concerned with the further claim that the most natural *terms* in which to respond to that metaphorical content exploit the initial metaphor as well. Again, I have already shown in §1.1.2.2 that hearers are not compelled to respond in those terms. But it is natural and felicitous for them to do so, and I need to explain this fact. So, Benvolio might naturally respond to Romeo's utterance of (56) by saying

(65) She sure is the sun,

or

(66) Juliet's not the sun: Rosaline is. *She's* the brightest, most beautiful girl around here.

If he does respond in this way, then he too speaks metaphorically: he takes up the very same aspect for his own use, and his utterance likewise depends for its construal upon the aspect.[59] My explanation of this fact is very simple. We naturally interpret Benvolio's utterance metaphorically because, just as with Romeo's utterance, a literal construal is patently false, and also because Romeo's utterance has already introduced the relevant characterizations and aspect into the conversation. These presuppositions remain in place unless someone explicitly disavows them.

Stern's discussion exploits more extreme versions of the same basic phenomenon of conversational presupposition. He points out that utterances of the following sentences are bad:

(67) The largest blob of gases in the solar system *is the sun*, and Juliet/Achilles *is, too*.
(68) Juliet *is the sun*, and Achilles *is, too*.

Because, he says, "it is difficult to see how we might account for these constraints [whose violation produces the badness] in terms of use or mutual beliefs and expectations," therefore "we must attribute to metaphor the semantic structure, or meaning, necessary for the requisite condition to apply."[60] I argued in §1.1.2.2 that these sentences *are* interpretable, and that the badness can't be a distinctively semantic product of the verb phrase ellipsis, as Stern claims, because the utterance is still bad when we split the conjuncts into independent sentences:

(67') The largest blob of gases in the solar system is the sun. And Juliet is the sun.
(68') Juliet is the sun. And Achilles is the sun.

I can now say how the constraints in question do in fact result from 'mutual beliefs and expectations'.

In (67), the description of the sun in literal, scientific terms renders a non-scientific characterization of the sun irrelevant to the conversation, and even counter-productive for the purposes of interpretation. But the second conjunct then exploits just such a non-scientific characterization as a means toward a metaphorical interpretation. Confronted with the second conjunct, the hearer must abruptly switch interpretive strategies — from literal to metaphorical — *and* radically reconceive his

---

[59] Recall from §4.5.3 that because the same 'essential description' characterizes both Juliet and to Rosaline, (66) involves just one aspect.
[60] *Metaphor in Context*, pp. 70-1.

operative characterization. Something similar also happens in (68), although here the incongruity between the two conjuncts is not as great. This is because, first, both conjuncts are metaphorical, and so require the same interpretive strategy. Second, in both cases the aspect is applied to a person, and so there is significant similarity in how the aspect applies. We do still encounter some serious incongruity, though, because Juliet and Achilles are thought of under very different 'essential descriptions'. As a result, many of the sun-features we used in applying the aspect to Juliet turn out to be irrelevant for applying the aspect to Achilles, and we need to hunt out new sun-features that can be matched to Achilles-features.

My own explanation of the infelicity here is actually not all that different from the positive explanation that Stern offers in chapter 4. He too invokes conversational presuppositions, and he too takes those presuppositions to be pragmatic, a function of conversational participants' 'mutual beliefs and expectations'. Thus, it is not so clear that his own view *requires* a semantic explanation, as he insists.

However, Stern also has a second sort of example, which I did not discuss in §1.1.2.2, which he also takes to show that interpretation is constrained in a distinctively semantic way, and in a way which motivates his specific account of metaphor. He claims that the features contributed by the metaphorical interpretation of 'is the sun' in

(69) Juliet might have been the sun

must be fixed relative to the *actual* context of utterance, not to an alternative possible context or circumstance.[61] That is, he claims, an utterance of (69) must be interpreted to mean that Juliet might have possessed those very features which she would be claimed to possess by a metaphorical utterance of the simpler sentence (56) in the same context of utterance: of being beautiful, good, nurturing, and so on. Stern takes this to show that metaphor behaves in a way precisely analogous to demonstrative terms, which also fix their values relative to actual contexts of utterance, regardless of the circumstances of evaluation. If this is right, then we do seem to have a genuinely semantic constraint on metaphorical interpretation.

---

[61] *Metaphor in Context*, p. 71.

Here, though, once again I think it is a failure of interpretive imagination which provides the constraint, not the nature of metaphor itself. Stern is imagining a context in which Benvolio[62] utters (69) in response to Romeo's initial utterance of (56), meaning that Juliet *is* not, but *might* have been, as Romeo characterizes her. Here, Benvolio uses his knowledge of Romeo's characterization of Juliet to generate the relevant aspect, and so to determine which features Juliet is being claimed to possess. Benvolio then claims that Juliet merely *might have* possessed *those* features. So far, so good.

However, it is also true that an utterance of 'is the sun' in (69), in what is effectively the same context of utterance, can be used to contribute the features that Juliet *would* have in a counterfactual circumstance of evaluation. Suppose, that is, that Paris, rather than Benvolio, utters (69) in a context in which Juliet and the sun have precisely the same features as they do in the case Stern is imagining. Paris says this, let us suppose, in the context of musing over what Juliet might have become had she not conceived this ill-fated passion for a Montague, but had married him instead to become a leading lady of Verona. In such a case, his utterance of (69) would primarily claim that Juliet might have been the most important socialite, the one to whom all other social figures deferred. But in this utterance of (69), 'the sun' contributes very different content than it did for Benvolio's utterance. For one thing, Paris's utterance makes no claims about the possibility that Juliet might nurture Romeo's soul, or even be morally good, as Benvolio's does. Benvolio's utterance treats Juliet's possible goodness and beauty as like the sun's light in being natural and original, while Paris's utterance predicates a highly socialized, artificial importance to Juliet.

On my view, the reason that Paris's utterance of 'the sun' in (69) contributes such different contents than Benvolio's does is that in interpreting Paris's utterance, the governing characterization of *being the sun* is applied to Paris's characterization of an imagined-alternative-Juliet, in the sense discussed in §4.7.4. By contrast, for Benvolio's utterance the aspect is applied to (Benvolio's understanding of) Romeo's characterization of the actual Juliet. As a result, in the two cases different

---

[62] Actually, Paris is Stern's favored interlocutor for Romeo, though this seems to be in tension with the plot of the play.

features get matched to different features in the same sun-characterization, or are matched to the same sun-features by different means.

Again, at first blush, Stern's positive explanation of metaphorical interpretation may not seem that different from mine. We both want to appeal to something like the same process of 'interaction' between background assumptions for determining the content of 'is the sun'. On his view, presuppositions about the subject in the actual context of utterance serve to 'filter out' otherwise contextually relevant presuppositions that are triggered by the predicate. The problem for Stern's theory comes when we turn to the specific mechanisms of his semantics. He adopts a Kaplanian theory of names and demonstratives, and so, because 'Juliet' is a proper name used literally, it functions as a directly referential term on his view. It thus contributes just the same object to the propositions expressed by both Paris's and Benvolio's utterances: Juliet herself. (Although of course *she* may have very different features in alternative circumstances of evaluation.) And on Stern's view, [Mthat 'the sun'] *also* functions as a directly referential term, to contribute to the proposition expressed just those features of the sun which are metaphorically relevant in the actual context of utterance.

The problem for Stern is that, in the imagined context of utterance in which Paris utters (69), Juliet as she is in that context 'filters out' the wrong presuppositions: Paris isn't interested in the properties which would be predicated of Juliet as she *actually* is, but as she *would* be. As a result, Stern lacks the semantic resources to access the properties he needs: those features of an imagined-alternative-Juliet which 'filter' in such a way as to deliver Paris's intended meaning.

I suspect Stern would respond to this objection by saying that I'm misconstruing the role that 'the actual context of utterance' plays in his theory. The two utterances of (69), by Benvolio and by Paris, do occur in different contexts. Sure, he might say, Juliet and the sun do by hypothesis possess just the same properties in the two contexts, and there is no indexical term like 'I' that could make the difference in speakers between the two utterances semantically relevant. But there are still important and relevant contextual differences. In particular, Benvolio's utterance occurs in response to Romeo's of the simpler, non-modal sentence

(56) Juliet is the sun,

which is itself embedded in a much larger utterance. And Romeo's utterance invokes a complex schema of objects (the moon, stars, a lamp, among other things), which contribute significantly to the relevant conversational presuppositions.[63] Paris's utterance does neither of these things.

I think this point is quite true and important, but it still doesn't address the problem I isolated with Stern's semantic story above. He still needs to appeal to a characterization of an imagined-alternative-Juliet in order to deliver the appropriate interpretation for 'might have been the sun', but his directly referential semantics prevents him from accessing this counterfactual characterization.

To address this problem, Stern must maintain the following position: although actual Juliet is contributed to the proposition expressed by Paris's hypothetical utterance, and although [Mthat 'the sun'] also functions in a directly referential manner to contribute the same properties to all circumstances of evaluation, still the properties which it contributes in this context of utterance to every circumstance of evaluation are fixed by a process of *merely pragmatic* interaction between [Mthat 'the sun'] and Juliet as she *would be* in the transformed circumstances in which Paris ultimately intends his utterance to be evaluated. The crucial point for Stern to insist on here is that the role that imagined-alternative-Juliet plays here is merely pragmatic and not semantic. The semantics of the [Mthat] operator call for *some* set of features to be contributed to the proposition prior to its evaluation in any particular circumstance, and the *same* features must be contributed by the context of utterance to *every* circumstance of evaluation. But in this case, just which set of features get contributed in that particular context of utterance must depend on accessing a particular circumstance of evaluation, which just happens to be contextually privileged in the context of utterance, in a merely pragmatic manner.

If this is not outright inconsistent, then at very least it renders Stern's claim that the content contributed by a metaphorical term is always fixed in the actual context of utterance, and so functions in a distinctively demonstrative-like way, quite precarious and unintuitive. On a less technical level, the fact that he needs to appeal in such a crucial way to a merely pragmatic interpretive process again renders his

---

[63] *Metaphor in Context*, pp. 156-9.

claim that metaphor is a semantic phenomenon less credible. A purely pragmatic explanation, like mine, can employ the same basic explanatory resources in a much simpler way. Thus, we do not need to postulate a special [Mthat] operator in order to account for the constraints on metaphorical interpretation. We need to appeal to "mutual beliefs and expectations" anyway, and they can do all the necessary work on their own.

5.4.6: *Metaphor as Make-Believe*

The last position on metaphorical comprehension that I want to take up here is one that I have not considered earlier. This is Kendall Walton's account of metaphor as "prop oriented make-believe";[64] it is relevant here in part because it may seem to be similar to my own view, and in part because David Hills suggests that it fixes the semantic content of metaphors. Walton sees metaphor as another manifestation of a phenomenon he takes to be central to understanding representation and aesthetics more generally: make-believe. On his view, games of make-believe are governed by "principles of generation"; these principles prescribe some imaginings, which are required in order to count as participating in the game at all, and make others merely permissible or optional. We often employ props, such as tree stumps, paper airplanes, and paintings, in these games. Usually these props are merely that: tools for facilitating the game itself. Sometimes, though, our participation is primarily guided by our factual interest in the prop, rather than by an interest in what is fictional for its own sake. Walton calls this sort of make-believe *prop* rather than *content* oriented.

Walton claims that in making a metaphorical utterance, a speaker "calls to mind" a game within which the utterance counts as a literal assertion about a prop in the game. Thus, Romeo calls to mind a game in which Juliet is a prop, and in which it is *fictional* that she is the sun — that is, in which we are to imagine that she is.[65] Sometimes the speaker merely suggests by her utterance that there is such a game; the game is just a tool for efficiently communicating her intended content. With more 'deeply meant' and

---

[64] "Metaphor and Prop Oriented Make-Believe."
[65] Of course, it is Shakespeare who makes it fictional that there is someone named Romeo who invites us to think of someone named Juliet in such a way that it is fictional (within the game of *Romeo and Juliet*) that she is the sun.

sustained metaphors, she makes the utterance as a participatory move within the implied game, and thereby draws her listeners into the game as well. In either case, by calling that game to mind, the speaker makes salient the *actual* circumstances that underwrite the fiction, and presumably claims *that* those circumstances obtain. Insofar as the facts that underwrite the fiction can be described directly, the game is disposable and the metaphor is paraphrasable, if often only in a laborious fashion. For some metaphors, the relevant facts may be accessible only from within the fiction, and so the metaphor will not be paraphrasable.

There are several ways in which Walton's (and in turn Hills's) account and mine are closely parallel. First, we both look to a rich and nuanced pattern of thought in which imagination plays an essential role, but which is still concerned with thinking about things in the world. We both find important commonalities between the sort of thought relevant to engaging with fiction and with metaphor, but also an important difference in orientation between the two. We both articulate a sort of genuine commitment to a proposition which is not commitment to its truth: where he invokes fictionality, I invoke fittingness. And we both think that for some metaphors the rich, nuanced and imaginative pattern of thought plays a merely instrumental role, while for others it is essential.

However, there are more serious differences between us as well. One difference is that Walton's analysis is entirely propositional: games of make-believe are cashed out as a series of propositions of the form 'it is fictional that *P*'. This includes propositions where *P* specifies our affective attitudes. By contrast, I think that a complex structure among propositions and affective attitudes (as opposed to merely pretended attitudes) is an essential component of our engagement with fiction, metaphor, and make-believe itself.

The most fundamental difference between us, though, concerns what we take to be the appropriate order of explanation. Walton takes games of make-believe themselves to be basic, and uses them to explain how we experience and think of things. I think these games themselves need to be explained in terms of our abilities and dispositions to see and think of things under aspects. Here is one illustrative example. Walton maintains that our ability to recognize the icons that identify rest rooms as

for men or women depend upon a game in which "one imagines recognizing a woman, or a man, or a person in a wheelchair; fictionally one does so."[66] But the postulation of a game of make-believe here seems quite otiose. Our ability to bring the relevant concepts to bear in perception is, I think, more basic, simpler, and itself explains why we engage in the games of make-believe we do, when we do engage in them. And in this case, I don't think we do engage in any make-believe: we just see the figure *as* a figure of a woman (or more accurately, as Wollheim says, we see a woman *in* the figure).

I think the same point goes for metaphor: our ability to characterize things and to think of them under aspects explains both our disposition to speak metaphorically, and our disposition to engage in certain games of make-believe, not the other way around.

Walton does admit that usually, in the cases he wants to explain by appealing to prop oriented make-believe, we don't actually participate in any make-believe: we merely *allude* to such a game, he says. He also admits that in many such cases, the putative principles of make-believe generation are obscure to the interlocutors. We are "cognizant" of them only in the sense that we have "abilities and dispositions" to respond in certain "appropriate" ways when confronted by pictures, rest room icons, and metaphors.[67] Given these points, in order to make sense of all the manifold cases in which we do apply a concept in perception or thought without full commitment to the truth of its applicability, Walton ends up postulating that our engagement with the world and each other is shot through with "hints of, allusions to potential and often fragmentary games."[68] But in all this, Walton has still not accounted for the psychological proclivities and patterns of thought that guide *which* principles of generation we adopt and allude to in which situations.

I think that once we do account for these proclivities and patterns of thought, we will have accounted for the 'abilities and dispositions' that Walton employs as evidence for make-believe directly, without postulating that we are almost always alluding to games that we're neither participating in nor conscious of. These patterns of thought will also go a long way toward explaining the explicit games of

---

[66] "Metaphor and Prop Oriented Make-Believe," p. 43.
[67] "Metaphor and Prop Oriented Make-Believe," p. 52.
[68] "Metaphor and Prop Oriented Make-Believe," p. 42.

make-believe in which we do engage — for instance, in reading fiction. I think, of course, that my account of characterizations and aspects is an account of precisely these patterns and processes of thought. I therefore think that in the end, Walton will need to appeal to something like my account of characterizations and aspectual thought in order to explain make-believe, while I don't need to appeal to his account to explain metaphorical comprehension. This lends additional credence to the claim that characterizations and aspects do independent explanatory work in a theory of mind.

5.5: *Conclusion*

Now that I have laid out my own view of metaphor and compared it with those of the most relevant alternative theories, I hope to have convinced you that my account can indeed do justice to the full range of metaphorical phenomena that we encountered in our survey of the three main types of theories in chapter 1, and can do justice to that phenomena in a unified and independently motivated way.

I have argued for a pragmatic theory of metaphor, on which a speaker says one thing in order to mean something other than (and sometimes, in addition to) what she says. What is distinctive about metaphor, as opposed to other pragmatic phenomena, is that the connection between what is said and what is meant is mediated by the cultivation of an aspect in thought. In the basic case, the relevant aspect (that is, the governing characterization) is given by the literal meaning of the verb phrase, and the subject to which the aspect is applied is given by the literal meaning of the noun phrase. In more complex cases, either the subject or the governing characterization must be inferred from contextual clues or from the (more or less conventional) colligates of the words uttered.

The content of the speaker's meaning then consists of predicating of the subject just those features of the subject characterization which are **most tightly matched** to the most prominent features of the governing characterization. Depending on the context and the linguistic resources of the language, this content may be more or less open-ended, more or less independently expressible in literal terms, and more or less difficult to determine. In addition, the speaker may merely be interested in the aspect as a

tool for isolating the features she wants to predicate of the subject, or she may be committed to the accuracy of the subject characterization as it is structured under the relevant aspect.

On the one hand, treating metaphor as a pragmatic phenomenon allows me to accept the fundamental intuition that in speaking metaphorically speakers manage to undertake the usual range of speech acts, without committing myself to a semantic theory of metaphor. More specifically, my account can do justice to the rich, complex connections between what is literally said and what is metaphorically meant without incorporating all that complexity within the semantic theory itself. This relieves the explanatory burden on the semantics considerably. By appealing to the associations that exist among particular words, as with colligation, and between particular words and particular characterizations, as with connotation, I can do justice to the intuition that metaphorical meaning depends in an important way on literal word meaning without committing myself to the unsustainably strong view that metaphor is exclusively an affair of words, as theorists like Levin, Beardsley and Kittay have maintained. And by appealing to the way in which pragmatic presuppositions alter hearers' interpretive expectations in a conversational context, I can explain the fact that led theorists like Recanati, Hills and Stern to a semantic, contextualist theory of metaphor: the fact that certain interpretations and responses are quite natural in response to a metaphorical utterance. But where they are forced to deny that alternative interpretations and responses are also available, I can explain this fact as well.

On the other hand, by locating aspectual thought as the means by which the content of those speech acts are determined, I can do justice to the rich cognitive and psychological effects that have led theorists like Davidson, Rorty, Moran, and Booth to claim that metaphor involves some essentially non-propositional component. By developing a relatively detailed theory of aspectual thought and the characterizations on which they are based, I hope to have removed some of the mystery from those non-propositional effects, to have connected metaphor to other related phenomena, such as aspectual perception and the interpretation of fiction, and finally to have pointed the way toward a richer, more encompassing theory of mind.

In chapter 6, I articulate more fully the lessons that I think our investigation of metaphor holds for a philosophical theory of language and of the mind. Before turning to this rather general topic, however, I take up the vexed question of whether metaphors can be paraphrased.

CHAPTER 6: METAPHOR, MEANING, AND THE MIND

In this chapter, I exploit my theory of metaphor, and the distinctions I have drawn in the course of developing it, to address two issues. First, can metaphors be paraphrased? Second, what can a philosophical theory of metaphor tell us about a philosophical theory of mind and of language? These questions are related at a deep level. On the one hand, the debate about whether metaphor can be paraphrased has been driven by common but largely unarticulated philosophical assumptions about what propositional content is, and about what sort of intentional structure needs to be in place if something is to be a candidate for paraphrase in the first place. In particular, philosophers have tended to assume a sharp division between rational cognitive processes — those processes which are governed by and answerable to rational law — and merely causal psychological processes — those more or less idiosyncratic processes which can only be the subject of empirical generalizations. Those facets of metaphor which fail to fit the usual rational model have been assumed for that reason also not to be fit candidates for paraphrase.

On the other hand, getting a detailed sense for how metaphor works, and how it is related to other linguistic and cognitive phenomena, provides us with a clearer sense for what would constitute an adequate theory of language and of mind. In particular, as I have said repeatedly, I think that close attention to the way metaphor works and the effects it produces itself calls into question the sharp division between the *merely* causal and the genuinely rational (which may be, and often is, causal as well). In turn, it also helps us to recognize that other, related patterns of thought and uses of language also fail to fit neatly onto just one side of this division. I turn to this claim in §6.2.

6.1: *Paraphrase and Its Discontents*

Too often, philosophical investigation into metaphor has gotten bogged down in the vexed debate about whether metaphors can be paraphrased, and so has failed to look hard at how metaphor itself really works. Disagreement about which metaphors are paradigmatic, about what even counts as a metaphor, and about what a metaphor means, all come to a head in this debate. Discussion is usually unproductive,

though, because no clear answers to these anterior questions have been provided. I start by taking up the positions on either extreme: that metaphors can never be paraphrased and that they always can. I then discuss what can be captured in a paraphrase and what is inevitably lost.

6.1.1: *No Content—No Paraphrase*

Recall that on Davidson's view of metaphor, speakers' metaphorical utterances merely serve to *cause* certain effects. These largely and most importantly consist in the non-propositional effects of 'seeing-as', although they may also include causing the hearer to entertain certain propositions. On a view of metaphor like Davidson's, there can be no paraphrase of a metaphor, simply because metaphors have no distinctive propositional content at all. Davidson writes:

> I agree with the view that metaphors cannot be paraphrased, but I think this is not because metaphors say something too novel for literal expression but because there is nothing there to paraphrase. Paraphrase, whether possible or not, is appropriate to what is *said*: we try, in paraphrase, to say it another way. But if I am right, a metaphor doesn't say anything beyond its literal meaning (nor does its maker say anything, in using the metaphor, beyond the literal). This is not, of course, to deny that a metaphor has a point, nor that the point can be brought out by using further words.[1]

The point of a metaphor, for Davidson, is of course to cause its hearer to see things in a new light: specifically, to nudge the hearer into noticing neglected likenesses, and perhaps also thereby to suggest certain attitudes and feelings toward the things being pointed out.

The first issue to address here is Davidson's claim that paraphrase is appropriate to — and only to — what is said. In a loose, intuitive sense of 'what is said', this is quite true: a paraphrase attempts to 'say the same thing in a different way'. It doesn't usually make sense to talk about paraphrasing things that aren't utterances, like trees or pictures. Paraphrase is appropriate only for propositional representations, and specifically only for those which are presented with some sort of illocutionary force. Paraphrase aims to specify the content and force of such utterances, but not other features like the tone of

---

[1] "What Metaphors Mean," p. 246.

voice or the language in which it is spoken.[2] We might thus think of paraphrase as a sort of indirect quotation, perhaps just *as* indirect quotation.

However, as we've repeatedly seen, we do use metaphorical utterances to communicate distinctive contents. Thus, for Davidson's claim that we *don't* say anything when speaking metaphorically to be even minimally plausible, he must be using 'say' in a more restricted, technical, quasi-Gricean sense. 'What is said' must isolate just what the speaker expresses simply in virtue of employing the words with their literal meanings, in the order they occur, in the context in which they are uttered, so as to determine a reference for any context-sensitive constituents the sentence may contain.[3]

The claim that Davidson is using 'say' here in this restricted technical sense, on which it is equivalent to what is semantically expressed by the sentence uttered, is supported by the fact that he is so concerned throughout "What Metaphors Mean" to argue against views of metaphor on which the *words* uttered themselves take on new meanings in a metaphorical utterance. Many of the theorists Davidson criticizes, such as Beardsley, Black, and Henle, failed (or refused) to distinguish clearly between semantics and pragmatics; and many people have accepted Davidson's arguments as defeating a semantic

---

[2] Focus and tone can sometimes alter an utterance's communicated content. This again forces the strict use of 'say' to depart from its intuitive use.

[3] One might well wonder whether Davidson is employing 'literal meaning' in the passage above in the usual sense of *conventional* meaning. On the one hand, that this is what he is doing is suggested by his claim that "Literal meaning and literal truth conditions can be assigned to words and sentences apart from particular contexts of use. This is why adverting to them has genuine explanatory power." ("What Metaphors Mean," p. 247.) On the other hand, this assumption would be in direct tension with his explicit rejection in "Nice Derangement of Epitaphs" of the idea that 'first meanings are governed by learned conventions or regularities" (p. 436).

Fortunately, I need not decide this exegetical matter here, because Davidson's claim in the passage in the text does not depend on whether he intends 'literal meaning' to be conventional or not. Suppose that throughout "What Metaphors Mean" Davidson only uses 'say' in the "Derangement" sense of 'first meaning'. Then my discussion below about how loose our usual use of 'say' is, and so how inclusive our usual practice of paraphrase is, will still go through. This is because replacing conventional meaning with 'first meaning' preserves — indeed, is designed to preserve — the distinction between literal sentence meaning and speaker's meaning. And the discussion below shows that our usual loose use of 'say' includes much of what must be thought of as speaker's meaning as opposed to literal meaning. Thus, 'first meaning' must still be significantly more restrictive than our usual loose use of 'say' is, if it is to do the theoretical work for which Davidson in "Nice Derangement" (following Grice while dropping Grice's assumption about conventionality) designed it.

We should also note, though, that Davidson's view in "What Metaphors Mean" may simply not be reconcilable with his later view in "Nice Derangement." In that later essay he explains Donnellan's speaker meaning as a case in which "Jones has said something true by using a sentence that is false. This is done intentionally all the time, for example in irony or metaphor" (p. 440). Here Davidson is himself employing 'say' in the loose sense to assert the very claim that he is so concerned to reject in the passage about paraphrase quoted in the text above and in "What Metaphors Mean" more generally. I can see no way of reconciling these claims. Thanks to John MacFarlane for pressing the point about conventionality.

theory of metaphor. This anti-semantic position is evident in a passage two paragraphs after the one just quoted:

> I depend on the distinction between what words mean and what they are used to do. I think metaphor belongs exclusively to the domain of use. It is something brought off by the imaginative employment of words and sentences and depends entirely on the ordinary meanings of those words and hence on the ordinary meanings of the sentences they comprise.[4]

But so far, this anti-semanticism leaves open — indeed, invites — the possibility of a pragmatic story, on which speakers can *mean* some other content by way of saying what they do. As we've seen, such a story is overwhelmingly plausible, given the uses to which we do regularly put metaphors. So, we now have three questions to answer: (1) Is Davidson right that paraphrase in general is appropriate *only* to 'what is said' in a narrow sense of that phrase? (2) Does Davidson maintain the strong claim that speakers can't even *mean* anything with determinate content by speaking metaphorically? (3) If so, is Davidson correct to maintain this? Davidson's claim that with metaphor there is no content to paraphrase can be true only if the correct answer is 'yes' to either (1) or both (2) and (3).

As to the **first question**: I think that what we're willing to report someone as having said does indeed track what counts as an appropriate paraphrase. But I also think that this applies to 'say' in the loose, intuitive sense rather than the technical one. Our ordinary use of 'say' is too loose and permissive to underwrite the technical Gricean contrast between 'what is said' and 'what is meant'.[5] The loose intuitive sense of 'what is said' is aimed at the content which someone *openly and obviously* commits themselves to by their utterance, and thus primarily draws a contrast with what someone merely insinuates or suggests.[6] To see that paraphrase covers more than simply what's said in the narrow technical sense, consider what we would normally take to be appropriate paraphrases of the following utterances:

(1) Can you pass the salt?

---

[4] "What Metaphors Mean," p. 247.

[5] Cf. Cappelen and LePore, "On an Alleged Connection between Indirect Speech and the Theory of Meaning."

[6] Thus, a paraphrase should not include non-obvious implications to which the speaker did not actively commit herself, such as the logical closure of what is explicitly said, even though she may in fact have committed herself to those implications by virtue of her utterance. It should probably also not include expressions of the same contents which the speaker herself would not be willing to avow if presented to her in the terms of the paraphrase, although this is much less clear.

(2) I was wondering whether maybe you'd be willing to consider writing a letter of recommendation for me.

In both cases, a natural, good paraphrase would explicitly report the request which the speaker is making only indirectly in her utterance: (1) and (2) should intuitively be paraphrased as:

(1') She's asking for the salt.
(2') She's asking for a letter of recommendation.

But these cases are paradigmatic examples of indirect speech or implicature. The speaker meant what she said, in the technical sense of 'said', but she was primarily concerned to make a further speech act. An appropriate paraphrase will report this. The technical Gricean distinction is essential for a clear, systematic theory of language use, but we should not pretend that it is transparently encoded in our normal ways of talking

Our intuitions are, I think, fuzzier about cases in which a speaker means something different from (as opposed to more than) what she technically says, but even here the standard of 'open and obvious commitment' seems to capture what is appropriately reported in a paraphrase. Thus, it would be highly misleading, though technically correct, to sincerely paraphrase a sarcastic utterance of

(3) What a brilliant idea: let's spend our last dollar on beer!

as

(3') She's saying that you had a smart idea which we should implement.

Rather, one should report the speaker as saying that this is not a good idea at all. (Note, though, that one can *sarcastically* paraphrase [3] as [3'].)

On the other hand, it is *not* appropriate to paraphrase what a speaker merely *insinuates* or suggests, at least not without explicitly flagging that this is what one is doing. If I respond to a question about how trustworthy John is by saying

(4) I think he's been out of trouble for a while now;

or to a question about how good a philosopher he is by saying

(5) He has had an interesting idea or two in the time I've known him,

or

(6) His command of English is excellent, and his attendance at tutorials has been regular,

then it seems that I should *not* be reported as having said that John is untrustworthy or not a good

philosopher, precisely because I have been careful precisely to avoid saying those things. I have clearly

chosen my words so as to refrain from openly and obviously committing myself to those claims.[7] I could

thus rightfully object to such a paraphrase by saying 'I didn't say *that*', whereas such an objection would

ring quite hollow for any of (1), (2), or (3) above.

==It should thus be clear that Davidson was wrong to tie paraphrase to the technical notion of 'what
is said', rather than to its intuitive loose use.== Now, as to our **second question**: does Davidson really

intend to claim that we don't say anything in speaking metaphorically beyond the literal — where 'say' is

now understood in the loose sense appropriate to paraphrase?

Although there is some scattered countervailing textual evidence, it seems pretty clear that

Davidson is committed to this claim: that a speaker doesn't *mean*, or actively commit herself to, any

determinate content in speaking metaphorically. He writes:

> Since in most cases what the metaphor prompts or inspires is not entirely, or even at all, recognition of
> some truth or fact, the attempt to give literal expression to the *content* of the metaphor is simply
> misguided.[8]

==Here the claim seems to be that there is usually *nothing* associated with a metaphor which could be
represented in a proposition, and therefore nothing which could be said by a metaphor, let alone said in
another way. As we've seen, this is simply false.== At the very least, metaphors often do prompt

recognition of heretofore neglected and even unnoticed facts, and these facts can in principle be

represented propositionally. Elsewhere, Davidson explicitly admits this, and lodges his objection to

paraphrasable content in a worry about the *connection* between what is said and the facts noticed:

> The common error is to fasten on the contents of the thoughts a metaphor provokes and to read these
> contents into the metaphor itself. No doubt metaphors often make us notice aspects of things we did not
> notice before; no doubt they bring surprising analogies and similarities to our attention... The issue does
> not lie here but in the question of how the metaphor is related to what it makes us see.[9]

---

[7] Notice, though, that (5) could also be used as a case of litotes or understatement (by, say, a British academic) to
deliver high praise. In this case, an appropriate paraphrase would report the intended content: that John *is* a good
philosopher.
[8] "What Metaphors Mean," p. 263.
[9] "What Metaphors Mean," p. 261.

and again, in a famous passage:

> Joke or dream or metaphor can, like a picture or a bump on the head, make us appreciate some fact, but not by standing for, or expressing, the fact. If this is right, what we attempt in 'paraphrasing' a metaphor cannot be to give its meaning, for that lies on the surface; rather we attempt to evoke what the metaphor brings to our attention. I can imagine someone granting this and shrugging it off as no more than an insistence on restraint in using the word 'meaning'. This would be wrong. The central error about metaphor is most easily attacked when it takes the form of a theory of metaphorical meaning, but behind that theory, and statable independently, is the thesis that associated with a metaphor is a definite cognitive content that its author wishes to convey and that the interpreter must grasp if he is to get the message.[10]

The position in these latter passages seems to be that even when what is noticed is primarily factual, the connection between the utterance and the noticed content is a merely causal one. If so, it cannot be treated in terms of what the *speaker* means. Rather, it can only be classified under what happens in and to the hearer. But as we saw in §1.3, this position too is false. Metaphorical utterances do serve to convey relatively distinct, more or less determinate contents. Further, they present those contents with illocutionary forces of the usual kinds. But this entails that the hearer has *understood* the speaker only if he understands what the speaker has committed herself to by her utterance.

The relation between what the speaker says, what she means, and what the hearer is supposed to grasp therefore fits the usual model of conversational interpretation perfectly well. More specifically, in speaking metaphorically, speakers don't just insinuate or suggest what they mean, as they often do in making implicatures. Rather, they openly and obviously commit themselves to propositional contents distinct from what they technically said. Thus, their intended metaphorical contents meet the test for appropriate paraphrase above.

Indeed, it looks like a metaphorical utterance will *always* have some content which is at least a candidate for paraphrase. (We still have to determine whether or not it always can in fact be paraphrased, and what a successful paraphrase will involve.) That is, if there are no matches at all between elements in the governing and subject characterizations, then the aspect just does not work, and there is no metaphor, only a literal utterance which is most likely absurdly false or trivially true. If, on the other hand, there are some such matches, then they determine at least some *M*-features, which in turn form the content of the

---

[10] "What Metaphors Mean," p. 262.

utterance. So, if Davidson is right to say that "there are no unsuccessful metaphors, just as there are no unfunny jokes," on the grounds that an utterance doesn't even count as a metaphor unless it succeeds in "bringing something off,"[11] then it is also true that there is always something for which a paraphrase can appropriately be attempted.

6.1.2: *Full Paraphrasability*

Taking these lessons to heart, it might now seem natural to conclude that we can *always* deliver a paraphrase of a metaphor's content. This is Searle's position. He maintains a Principle of Expressibility, on which "whatever can be meant can be said."[12] Thus, if something is the content of a speaker's meaning, then it can in principle be expressed literally. And if it's not content, but mere "psychological effects," then a theory of meaning has no business with those effects anyway. We need to address two issues here. (1) What does Searle mean by the claim that "whatever can be meant can be said"? (2) How does this claim apply to metaphor?

Searle offers the following as a more explicit statement of the Principle of Expressibility:

> for any meaning *X* and any speaker *S* whenever *S* means (intends to convey, wishes to communicate in an utterance, etc.) *X* then it is possible that there is some expression *E* such that *E* is an exact expression of or formulation of *X*.[13]

I think we can take it that by **'meaning'** Searle means truth-conditional, propositional content; he explicitly distinguishes between meaning and mere 'effects' immediately after the definition above. What we need to understand is what he means by 'it is possible that': what sort of possibility is this, and what are the constraints on it?

Some theorists of meaning might maintain that a language must contain expressions sufficient for the literal expression of any meaning a speaker may want to convey, on the Tractarian grounds that "the limits of my language" are "the limits of my world."[14] Searle, however, cannot maintain this view, not

---

[11] "What Metaphors Mean," p. 245.
[12] *Speech Acts*, p. 17.
[13] *Speech Acts*, p. 20.
[14] *Tractatus Logico-Philosophicus*, §5.6.

least because he takes linguistic representation to derive from the mind's Intentionality, and not the other way around. *We* make words mean what they do because of the thoughts *we* have. For this reason, he explicitly admits that a language "may not contain the words or other devices for saying what I mean."[15] In this sense, it may *not* be possible, given the semantic resources of the language, to find a literal and exact expression of one's thought. Indeed, Searle points to the need for metaphor as proof of just this point.[16] Rather, he interprets the Principle of Expressibility so that when there *is* a semantic inadequacy of this sort, then "I can in principle at least enrich the language by introducing new terms or other devices into it."[17]

So we now need to determine whether such '**enrichment**' is in fact always in principle possible. It seems that the relevant 'enrichment' can be accomplished either by definition or by ostension. *Explicit* definition cannot itself enrich the semantic resources of the language, because the new term is being introduced precisely as equivalent to some existing, more complex expression. This is not true in the case of *implicit* definition; but it is still true that in order for a term to be introduced by implicit definition, there must be some set of sentences which jointly determine what the undefined term must refer to if those sentences are to come out true. At the very least, there must be some common, canonical practice of using the term which likewise fixes a referent. But this means that there must be some common, if merely implicit and perhaps inchoate, representation of the thing or property referred to: as whatever it is that meets all the requirements imposed by the use of the term in the relevant set of sentences or linguistic practice. A speaker will then be able to exploit this common representation to introduce the new term.

When we turn to ostension, however, this situation — that there be an existing, public representation on the basis of which the new term can be introduced — does not always obtain. As a result, a speaker may not always be able to enrich the common language so that she can express just what she means. As we saw in §5.3.2, a speaker may sometimes have had a causal interaction which enables her to think about some property *P* which is not denoted by any standing term in the language. She

---

[15] *Speech Acts*, p. 19.
[16] "Metaphor," p. 114.
[17] *Speech Acts*, p. 20.

herself will be able to refer to that property in demonstrative terms, by exploiting that causal connection. In the cases we discussed in §5.3.2, she may be able to refer to the particular gesture made by a ballerina, or the particular feeling aroused by listening to the Fauré Requiem, by saying something like 'that gesture the ballerina made' or 'that feeling I had then.' But, unless that property is present in the common environment shared by the speaker and her hearers, or unless her hearers at least share the same means of ostending the property through their memories, her hearers will not be able to grasp what the referent of her demonstrative is. In such a situation, the speaker cannot introduce a literal term for the property she is thinking about into the language, because she and her hearer have no common access to the property in question. However, metaphorical utterances like

(7) She was a swan skimming over a smooth lake

or

(8) The Fauré Requiem puts you on the top of a lonely, icy mountain crag

may succeed in communicating what the relevant property is. They do this by exploiting the role that an analogous but distinct property plays in our characterization of another object in order to introduce the intended property into our characterization of the subject under discussion.

Searle is aware of this possibility, but denies that it poses a relevant objection, by denying that the Principle of Expressibility should apply to expressibility in a *public* language:

> The principle that whatever can be meant can be said does not imply that whatever can be said can be understood by others; for that would exclude the possibility of a private language... Such languages may indeed be logically impossible, but I shall not attempt to decide that question in the course of the present investigation.

I do not think this is a plausible interpretation of the Principle of Expressibility. If we accept the qualification that Searle proposes here, then the Principle will be so weakened that it will no longer be capable of doing significant theoretical work. In particular, it will no longer be relevant to the discussion at hand, of whether a metaphor's propositional content can always in principle be conveyed in other, literal terms. This discussion turns on whether literal means are always in principle available for making one's metaphorical content understood. Indeed, Searle himself draws as a consequence of the Principle of Expressibility the more general claim that "nonliteralness, vagueness, ambiguity, and incompleteness are

not theoretically essential to linguistic *communication*."[18] But metaphor and nonliteralness more generally are theoretically inessential to linguistic communication only if other means of communication are always in principle available. If the Principle of Expressibility is weakened so that it no longer requires that what can be meant can be said in terms that others can at least in principle understand, then the Principle obviously cannot be used to derive this claim.

Of course, speakers will not always succeed in getting their hearers to understand them in particular cases: either the speaker or the hearer, or both, may not be sufficiently resourceful. The speaker may not hit upon a metaphor that generates an appropriate match; the hearer may not figure out just how the new, intended feature is supposed to be matched to the familiar feature in the governing characterization, and so not be able to identify what the intended feature *is*. But in such a situation, communicative success is still in principle possible.

In sum, the question we are investigating is whether speakers can always in principle introduce terms into a language that will enable them to say what they mean in literal terms. Given the possibility articulated in §5.3.2, that the speaker has had past causal interaction with a type of object or especially a property for which there is no literal term in the language, and that the hearer lacks the same sort of causal access to that object or property, the answer to that question must be 'no'. A speaker may not always be able to find literal means, even in demonstrative terms, to articulate the content she has in mind in such a way that her hearer can in principle understand it. Therefore, given that speakers turn to metaphor precisely in order to talk about these sorts of objects and properties, the content of at least some metaphors cannot in principle be fully paraphrased.[19] Perhaps, *after* having gotten the hearer to identify the appropriate property by metaphorical means, the speaker can then introduce a new, literal term with this very feature as its denotation. This certainly does happen; it is part of what makes metaphor an

---

[18] *Speech Acts*, emphasis added.

[19] One might also be driven to speak metaphorically in order to claim that the subject possesses the higher-order property of a property or an affective response being *fitting* for the subject. In these cases we also lack an existing alternative literal means of specifying the property in question. However, if my story about fittingness is correct, then we should now be able to exploit fittingness-talk in paraphrasing such metaphors (see below). Perhaps the most compelling cases here are of it being fitting for one to experience an emotion or other attitude in relation to the subject; think here of trying to convey the appropriate response to an abstract modernist artwork, such as a Rothko.

important tool in the history of a language. However, I take it that in such a case the metaphor itself plays an essential role in defining the new term, and so that this possibility proves neither that a genuine alternative, literal means of expression is in fact available, nor that metaphor is not theoretically essential to linguistic communication.

6.1.3: *Possible Partial Paraphrases*

So far, we've seen that both extreme positions on the paraphrasability of metaphor are misguided: Davidson is wrong that there is nothing there of the right sort to be paraphrased, and Searle is wrong that what content *is* there can always in principle be paraphrased. Now I want to explore the middle ground, paying attention to the variables among metaphors we discussed in §5.3. Can at least some metaphors be fully paraphrased? Do others metaphors remain beyond the reach of full paraphrases, for reasons other than failure of independent semantic communicability?

A full paraphrase can be given of metaphors which occur in minimally weighted conversational contexts, and whose content therefore comprises only a few features, just in case that content consists of objects and properties denoted by alternative semantic expressions that the hearer can in principle understand. Routinized metaphors like

(9) Sam is a pig.
(10) Clifford is a mouse.
(11) Sally is a block of ice.

have a relatively determinate content of this sort. As we move away from heavily routinized metaphors, even if we remain within relatively unweighted conversational contexts, the content of the speaker's meaning tends to become less determinate and more open-ended. For instance, in the exchange

(12) Alice: How's work going?
    Bill: Still on the chain gang,

Bill's metaphor does not rely on a particularly novel or distant aspect, and is not intended to induce a particularly rich pattern of thought in Alice. Still, there is significant indeterminacy in its content. It conveys some complex of claims in the following rough range: Bill doesn't enjoy his job, he finds it

routine and constricting (probably both emotionally and intellectually), he feels the management is more or less authoritarian, but he also doesn't feel like he can quit.

Now, does this specification adequately paraphrase Bill's intended meaning? We might think not: if we presented Bill with this paraphrase, and asked whether he meant precisely *those* claims, Bill would likely respond that he meant more or less something in that ballpark, but hadn't sorted out the particular claims among their rough alternatives. His thought was itself relatively rough. This is an important point, and one that philosophical theories of language have tended to neglect in the past. However, as we saw in §1.3, such indeterminacy and open-endedness is not distinctive of metaphor. Much, perhaps even most of our talk is 'loose' in this way.[20] Paraphrases of these literal utterances will therefore likewise be partial and approximate, and will impose a determinacy on the speaker's meaning that it did not initially possess. One of the great benefits of loose talk, and of metaphorical speech specifically, is that it enables us to communicate the contents of our thoughts in a mode that retains the rough, somewhat indeterminate, but complex form in which we often actually think. Speaking clearly can be so difficult precisely because it forces us to refine our thoughts themselves, and not just to find an accurate expression of them.

A similar point goes for the role that affective attitudes play in metaphorical communication. As we have seen, we often turn to metaphor because we want to communicate a particular affective attitude, and perhaps in particular to express that it is fitting for the subject under discussion. So, for instance, Raskolnikov's utterance in *Crime and Punishment* of

(13) I've only killed a louse, Sonia

serves to express his contempt toward the old money-lender, and specifically to make his loathing of her seem fitting, and thereby justified. We might think that precisely because affective attitudes are non-propositional, the role they play in metaphor must be incapable of being captured in literal terms. However, this obviously need not be so. As I said in §1.3, many literal terms also convey affective

---

[20] Cf. Sperber and Wilson, "Loose Talk"; Bach, "Speaking Loosely"; Dennett, "True Believers," "Brain Writing and Mind Reading."

attitudes, whether as part of their semantics itself or of their associated connotation. So, Raskolnikov himself goes on to paraphrase his utterance thus:

(13') I've only killed a louse, Sonia — a useless, loathsome, harmful creature.

The later adjectives make explicit several of the initial metaphor's underlying matches. And they do so in literal terms that, like the metaphor itself, attempt to justify his claim that the old money-lender is despicable, and so that his attitude of contempt, and even his murder of her, are warranted. Now, not all metaphorical utterances will be like this. As I said in §3.5.2, emotions, especially those aroused in response to what Novitz calls "certain complex and demanding situations"[21] are the sorts of things for which we most often lack alternative, literal means of expression. As a result, metaphors that express them are most likely to be unparaphrasable. But when this is so, this is a fact about the expressive resources of the language, not about the role of emotions in metaphor *per se*.

Someone might argue on more general grounds that because emotions and affective attitudes are not propositional and register no truth-conditional difference, they don't belong in a paraphrase in any case. At the very least, the fact that this argument would also apply to many literal utterances blocks any attempt to show that metaphors are thereby *more* essentially non-propositional and unparaphrasable than literal speech. I think further that the fact that we often maintain that certain emotions are *fitting* in response to certain individuals and situations means that there is a form of genuine commitment here after all. An adequate paraphrase, whether of literal or metaphorical speech, should capture the commitments which the speaker intends to undertake by means of her utterance. A theory of meaning which systematically excludes all expressive content on the grounds that it is not truth-conditional will be too narrow to accommodate this.

So far, I have surveyed two challenges to the claim that metaphors can be paraphrased — the roughness or indeterminacy of metaphorical meaning, and the role that affective attitudes play in metaphor — and have concluded that at the very least, these challenges apply to literal as well as metaphorical utterance. A **third** challenge does seem to be unique to metaphor. The more novel and

---

[21] *Knowledge, Fiction, and Imagination*, p. 120.

distant the aspect is on which a metaphor relies, the more likely it is that the most natural paraphrase will itself employ metaphorical language. The paraphrase will indeed spell out the subject's *M*-features, but in the terms of the same aspect. For instance, Cavell's paraphrase of

(14) Juliet is the sun,

which I first quoted in §4.5.3, is itself largely metaphorical:

> Now suppose I am asked what someone means who says 'Juliet is the sun'....I may say something like: Romeo means that Juliet is the warmth of his world; that his day begins with her; that only in her nourishment can he grow. And his declaration suggests that the moon, which other lovers use as emblems of their love, is merely her reflected light, and dead in comparison; and so on.[22]

Some theorists have taken this fact to show that a speaker's intended metaphorical content can't be fully paraphrased: we remain fixed within a **metaphorical circle**. However, this conclusion obviously doesn't follow. First, metaphorical paraphrases of metaphors seem natural and appropriate because they make manifest the matches on the basis of which the features in the subject characterization are identified. But this basis for matching is not essential to the metaphorical content itself. (I return to this point in the next section.) **Second**, we saw in §4.3 that aspect-application is often recursive. Matches between features within an aspect may themselves depend on metaphorical similarity, but those metaphorical similarities are always themselves both less complex and more familiar than the initial aspect. So too here: Cavell's paraphrase, like all competent paraphrases, is more explicit and less metaphorical than the original utterance. If we like, we can go on to paraphrase the paraphrase's constituent clauses, and *their* paraphrases will in turn be less metaphorical. In most cases, we will eventually 'bottom out' into non-metaphorical matches. In some cases, a basic match may still be metaphorical, but if so, the metaphorical similarity in question will be so fundamental and familiar, like the similarity between cold objects and cold personalities, that the metaphoricity is either highly obvious, or else nearly counts as a case of polysemy.

---

[22] "Aesthetic Problems of Modern Philosophy," pp. 78-9.

Now, most of those who decry "the heresy of paraphrase"[23] would accept most of the claims I've made about paraphrase so far. They simply want to insist that a paraphrase is inherently merely partial and approximate: it can never do justice to or replace the original metaphor. In this vein, Max Black claims that some cognitive content must in principle be left out by a paraphrase:

> Up to a point, we may succeed in stating a number of the relevant relations between the two subjects [the 'topic' and the 'vehicle']. But the set of literal statements so obtained will not have the same power to inform and enlighten as the original. For one thing, the implications, previously left for a suitable reader to educe for himself, with a nice feeling for their relative priorities and degrees of importance, are now presented explicitly as though having equal weight. The literal paraphrase inevitably says too much — and with the wrong emphasis. One of the points I most wish to stress is that the loss in such cases is a loss in cognitive content...[24]

Black makes two distinct, though related, complaints here. The first is that a paraphrase "says too much," by rendering explicit and determinate what was before implicit and approximate. I've already pointed out, in regard to this complaint, that a paraphrase of many literal utterances would err in just the same way. But if this is an inevitable effect of paraphrase generally, then it does not demonstrate any distinctive richness or ineffability for metaphor.

The second complaint is that a paraphrase always says what it does "with the wrong emphasis": it 'levels out' the structure of what we understand when we understand a metaphor. This is the more important point for decriers of paraphrase, and it is most pressing for rich, resonant, 'deeply meant' metaphors. As Cleanth Brooks says, a paraphrase "is not a rack on which the stuff of the poem is hung; it does not represent the 'inner' structure or the 'essential' structure or the 'real' structure of the poem."[25] It should by now be obvious that I wholeheartedly agree with Black and Brooks here: a complex structure among thoughts plays an absolutely central role in the patterns of thought that motivate us to speak metaphorically and which we grasp when we understand a metaphor. One of the most vital ways in which metaphors serve as effective communicative tools is that they allow us to communicate whole characterizations, with all their complexity and structure, and with their intimate intertwining of factual and affective modes of representation.

---

[23] Cleanth Brooks, "The Heresy of Paraphrase."
[24] "Metaphor," p. 46.
[25] "The Heresy of Paraphrase," p. 182.

The fact that this is an important feature of metaphors, though, is not (or need not be) the real bone of contention between defenders of paraphrase, like Searle, and its decriers, like Black and Brooks. What is crucially at issue is whether the complex structure is part of a metaphorical utterance's intended *content*. Black explicitly insists in the passage above that it is, while Searle would deny this. And here I think we encounter an important ambiguity, or at least unclarity, in the use of the term 'content'. In chapters 2 through 4, I tried to articulate the truth behind the intuition that there is something importantly non-propositional about aspectual thought — a way in which aspectual thought does not fit the sentential model of thought. If 'content' is understood as being elliptical for or necessarily equivalent to 'propositional content', then by definition neither structure nor affective attitudes can count as content.

If, on the other hand, 'content' is understood as elliptical for 'representational content' — for how the world is represented as being — then, I have argued, the structure and affective attitudes of characterizations *can* be part of a metaphor's intended content after all. Characterizations that differ only in the structural roles of prominence, centrality, and depth which they assign to their constituent features still differ significantly in how they represent their subjects as being. Likewise, characterization that differ only in which affective attitudes they represent as fitting for the subject still differ significantly.

For most metaphorical utterances, the structure of the subject characterization under the aspect forms no part of what the speaker intends to claim or otherwise commit herself to by her utterance. In these cases, I said, the aspect plays the merely instrumental role of getting the hearer to isolate the appropriate *M*-features. Even quite rich, resonant metaphors often depend upon merely instrumental aspects. But in §5.3.4, I also argued that for some, 'deeply meant' metaphors, the speaker *is* committed to making the claim that the subject characterization as it appears under that aspect accurately represents the subject's true character. These are the cases I think Black has in mind. In such cases, the characterization, with all its structure, *is* part of the metaphor's intended representational content.

Now, can this 'content' too be paraphrased? Our ordinary language contains no easy resources for talking about these structural relations, and so a paraphrase will be difficult to come by. But it is within our ability to develop ways of expressing these relations. In particular, if we exploit the account of

aspectual thought I have developed, then we should be able to preserve this structure within our paraphrase, by specifying the relative prominence, centrality, and depth of the content's constituents. We might do this in a relatively convenient manner by assigning relative numerical weights to inferential and characterizing relations, or by drawing visual networks to represent them, as I have.

Once again, this would involve quite a lot of effort, which would itself force us to refine our intended metaphorical content significantly. This is especially true when we are concerned with the characterization's structure. As I have stressed at several points, although structure plays an essential role in aspectual thought, our sense of it is normally quite rough and intuitive. We don't explicitly represent our characterizations, with their structure, to ourselves. Rather, we employ them in the course of representing things in the world around us. Again too, however, none of these points are essential to metaphor. Speakers often intend for their literal utterances to suggest an complex structure of implications, whether a characterization or not. The most accessible paraphrase of these utterances will likewise 'level out' that structure, and a more accurate paraphrase will take a lot of work to develop. It is true that we speak metaphorically especially when we are focused on communicating something about the structure of those implications. This is because metaphors *require* aspectual thought, and so the complex structure of characterizations, in order to determine speakers' primary communicative intentions, as other utterances don't. But again, there need be nothing mysteriously and distinctively ineffable about metaphor here.

With these points in mind, we can now answer a worry of Davidson's about the passage by Black which I quoted earlier. Davidson objects:

> If a metaphor has a special cognitive content, why should it be so difficult or impossible to set it out?... Why does Black think a literal paraphrase 'inevitably says too much — and with the wrong emphasis'? Why inevitably? Can't we, if we are clever enough, come as close as we please?...It should make us suspect the theory [that metaphors have a cognitive content] that it is so hard to decide, even in the case of the simplest metaphors, exactly what the content is supposed to be.[26]

Davidson is right to reject Black's claim that a literal paraphrase is inevitably wrong, but he is right for precisely the wrong reason. Davidson, of course, thinks the reason is that there is nothing there to be

---

[26] "What Metaphors Mean," pp. 260-2.

paraphrased. We can now say instead that the reason Black is wrong is that we *can* often "come as close as we please" to an adequate paraphrase — assuming we do have alternative semantic expressions for the intended *M*-features. Doing so will take a lot of work, but this shouldn't make us suspect that metaphors don't express any content, as Davidson thinks. Especially for the simplest and most routinized metaphors, but more generally as well, we *do* have a pretty good sense of that content. And in both the simple and complex cases, the sort of indeterminacy exhibited is not distinctive of metaphor, but is typical of speaker's meaning more generally. It is also true that for some metaphorical utterances, we care a lot about exploring and refining that indeterminacy, about articulating just which features the speaker must have intended, and in just what structure. But this is often because those metaphorical utterances are objects of distinctively aesthetic or theoretical interest.[27] This means that the details of the intended features and how they hang together matter to us in a way they don't for ordinary 'loose' talk.

I hope that this discussion has cast new light on the intuition that metaphors are essentially 'non-propositional'. In chapter 2, I isolated several ways in which both perception and metaphorical meaning differ from typical sentential forms of representation. ==These were first, that the content of perception and of metaphors is much more precise and rich than that of most sentences; second, that perception and metaphorical understanding exhibit a complex structure; and third, that moods and other affective attitudes can play a permeant, structuring role in aspectual perception and thought.== We can now see that, with one possible exception, these distinctive features of perception and metaphor are not actually incompatible with propositional forms of representation. Rather, they are simply absent from the content of many — though by no means all — literal utterances.

As to the **first** point: demonstrative terms can contribute to the contents of sentences representations which are just as precise as our powers of perceptual discrimination allow; and massively long strings of sentences can accommodate as much richness as we please. As to the **second** point: we can, if we choose, formulate higher-order propositional representations which specify the structural and fitting relations among the various constituent propositions of characterizations. Only the **third** point

---

[27] See Hills, "Truth and Aptness in Verbal Metaphor."

remains as truly non-propositional: affective attitudes represent their objects in substantive ways which cannot be captured by specifying any conditions which, if true, would make the attitude count as veridical or as otherwise satisfied. But once again, this genuinely non-propositional component of metaphorical understanding is not confined to metaphor, but is present in communication and thought more generally. When a thinker is committed to the emotion in question being fitting, this sort of commitment deserves to be included within a complete theory of meaning.

The initial intuition about the 'non-propositionality' of perception and of metaphor is thus best put, not as a claim about an essential difference in kinds of content, but as the claim that the most *natural and efficient* way of representing the content of perception or of metaphor does not consist entirely in a list of sentences that contain no demonstrative terms or terms with strong affective connotations.

6.1.4: *What is Unparaphrasable in Metaphor*

So far, I have said that although finding an appropriate, accurate paraphrase of a metaphor's intended content is often exceedingly difficult, it is often also theoretically possible. The only in-principle obstacle to paraphrase is lack of common semantic access to the features the speaker wants to predicate of the subject. However, one important feature of metaphorical utterances *will* inevitably be left out of even the most arduously detailed paraphrase. This is aspectual thought itself, and in particular the experience of twofoldness to which it gives rise. A paraphrase can in principle capture the speaker's intended *subject* characterization, with just the constituent features structured in just the way they would be if it were being governed by the aspect. It cannot capture the aspect itself, however: that is, the way in which the governing characterization applies to and structures the subject characterization. One way to put this point is to say that a paraphrase aims to 'say the same thing in a different way', and that the aspect in this case just is the particular 'way' in which one thinks about, and speaks of, the subject and its

possessed features. To find a different way of saying or thinking about how the subject is in the case of metaphor just *is* to find some mode of expression or thought other than the aspect.[28]

Another, more accurate way to put the point that a paraphrase cannot include the aspect itself is that the aspect forms no part of the satisfaction-conditions of the speaker's utterance. Suppose, for instance, that Benvolio and Romeo come to agree that Juliet has just the character that thinking of her as the sun would leads us to expect: she has just *those* qualities, structured in just *that* way, and just *these* emotional responses to her are fitting. Suppose, in fact, that their shared characterization of Juliet accurately represents her real character. However, suppose that they disagree about how the sun itself is, and that both of their characterizations of the sun are wildly inaccurate. In such a case, I think that Benvolio would still be in agreement with Romeo's intended claim, and that Romeo's intended claim would still count as true. It would, of course, a bit of a miracle that Romeo managed to express this truth by these inappropriate means, and that Benvolio managed to understand him. Thus, someone might rightly criticize Romeo's utterance for not being an apt means of expression, perhaps even for being infelicitous. But I do not think this criticism would undermine the truth of Romeo's utterance. If not, the aspect itself cannot be part of what a speaker commits herself to by her metaphorical utterance.

Searle puts a similar point to my claim that an accurate paraphrase will not include the aspect itself by saying that something must necessarily be lost in a paraphrase, because "without using the metaphorical expression we will not reproduce the semantic content which occurred in the hearer's comprehension of the utterance."[29] This fits with my claim that cultivating the aspect typically serves as the means, rather than the sole end, of metaphorical comprehension. For Searle, none of the aspect's 'non-propositional' effects are ever part of what the speaker means, and so they are not themselves part of what should be paraphrased. As I've said, I disagree with Searle on this latter point. I want to claim that the overall subject characterization, structured in just the way it is under the aspect, *can* sometimes be part

---

[28] All utterances involve expressing content in a particular way, and any other 'way' of expressing that content is likely to change what is communicated at least slightly. What is distinctive about metaphor is that the 'way' is through aspectual thought.

[29] "Metaphor," p. 114.

of what the speaker intends to communicate and to commit herself to as well. That is, the overall subject characterization can be part of the speaker's intended end and not just her means.

However, the same point does not go for the governing characterization; here Searle's point is fully correct. I think further that Searle's point reveals, not just an incidental fact about metaphor, but the explanation of an important intuition about it. In §4.7.2, I emphasized that one reason the perceptual metaphor of 'seeing-as' has seemed so natural for talking about aspectual thought is that aspectual thought, like the perception of pictures, involves a kind of 'doubleness': one sees or thinks about a complex representation *through* seeing or thinking about something else. For aspectual thought, this means two things. First, the governing characterization provides an overall organizing principle for structuring the subject characterization. Second, even after the governing characterization has molded the subject characterization in its image as fully as possible, the aspect continues to tug at that subject characterization, by reminding us of features which remain unmatched, and of the different structural roles that each element in a pair of matched features continues to play in their respective characterizations. Precisely in order to provide an accurate representation of the subject characterization itself, a paraphrase must ignore these remaining **differences** between the two characterizations.

The fact that an accurate paraphrase must necessarily leave out the aspect *per se* has several important ramifications, which help in turn to explain why we care so much about the aspect. First, we often derive a distinctive sort of aesthetic pleasure from seeing how an aspect applies: we enjoy discovering the rich connections between two very different characterizations. Second, the aspect provides a singularly efficient means for representing the speaker's intended content. But third, and more importantly, by applying the aspect, a hearer not only comes to understand the speaker's subject characterization as it currently stands. He also has a better sense for 'how to go on' in molding that characterization as new facts are encountered. He acquires, not just a static representation of the speaker's mental state, but a dynamic principle for thinking the way the speaker does. This is especially important in poetic contexts, in which the speaker cedes some of the responsibility for fixing the details of the content to the hearer. The hearer may also thereby gain a principle for applying related aspects to

other, related subjects as well. Thus, for example, by describing Juliet as the sun, Romeo also implicitly suggests at least a rough characterization of Rosaline, of love and its varieties, and of the natural world's relationship to the sun. In short, the aspect, in its application to a specific subject, can also serve to suggest a more general outlook in a way that a detailed elaboration of its consequences cannot.

6.2: *Rationality, Causality, and the Mind*

The question about whether metaphors can be paraphrased could only really be answered once we had a unified theory of metaphor: one which explained both metaphor's propositional, communicative content and also its 'non-propositional' effects, and which could therefore do justice to both the everyday conversational uses of metaphor and also its more rich and nuanced uses in poetic contexts. Only then could we get a sufficiently broad overview of the phenomenon of metaphor, and so deliver fully general answers about what metaphor can and cannot accomplish. I now want to step back even further, to explore what my investigation of metaphor can show us about language and the mind more generally, and what lessons it holds for philosophical theorizing about these matters.

If asked how metaphor works, the most common response by philosophers in general and by philosophers of language in particular would likely be to agree with Grice, Searle, and me that in speaking metaphorically, speakers exploit conventional semantic rules in order to communicate distinct contents. But if they were then asked how speakers manage to do *this*, I suspect that many philosophers would respond that this is a very interesting topic for psychological and sociological investigation, but not a distinctively philosophical matter, at least not beyond a general appeal to something like 'contextually salient respects of resemblance'. As Frege said, philosophy concerns itself with psychology only so far as to discern the laws of pure reason, the laws which determine how we *ought* to think. It does not attempt to discern from an armchair the general patterns of how we actually *do* think, and in particular what particular sorts of associations we happen to make between objects in particular sorts of situations.[30] The common philosophical assumption seems to be that anything beyond the general outline of an account of

---

[30] "Thoughts."

metaphor will inevitably delve into just these sorts of factual particularities. Because of these assumptions, many mainstream philosophers of language and especially of mind have not felt the need to think particularly hard about metaphor. It is enough, they think, to locate metaphor within the overall topography of linguistic phenomena, to indicate in broad outline how metaphor works, and then to turn the investigation of the empirical details over to other disciplines.

There have, of course, been some notable exceptions to this general view: my dissertation has discussed these very philosophers. But even they have tended to share the view that metaphor is a phenomenon of relatively marginal philosophical interest. Thus, the bulk of Davidson's interest in metaphor seems to be to demonstrate that it has no place within a theory of meaning, because it is not a case of meaning at all. Grice confines himself to explaining metaphor in terms of a vague appeal to 'resemblance in some respect or other'; Bach and Harnish largely follow his lead here. Searle has thought harder about the mechanisms of metaphorical communication than most philosophers, but as we have seen, he too offers not much more than a sketch of a theory, and concludes that metaphor is "not theoretically essential" to a theory of communication.[31] More recently, philosophers like Kittay, Fogelin, Moran, Stern, and Hills have paid more sustained attention to metaphor, but few philosophers outside of the narrow bounds of 'metaphor studies' have paid much attention to this work.

Throughout the last five chapters, I have tried to demonstrate that this benevolent neglect of metaphor is not a satisfactory position. First, at a very minimum, we have not conclusively shown that the basic claim about metaphor above — that in speaking metaphorically speakers exploit conventional semantic rules to novel pragmatic ends — is correct until we have constructed a complete and workable theory of metaphor on its basis. Until we have explained how all the distinctive phenomena of metaphor arise in terms of a coherent pragmatic theory, we have not demonstrated that metaphor is not in fact a semantic or non-cognitive phenomenon. We have merely asserted that it must somehow be treatable in pragmatic terms. And if metaphor were to turn out to be a non-cognitive or especially a semantic

---

[31] Black and Goodman offer more extensive consideration of the cognitive mechanisms on which metaphor relies, but their more general commitments about language prevent them from arriving at fully satisfactory views on the matter.

phenomenon, then our overall map of the topography of language use would need to be modified significantly. Philosophers like Stern and Hills claim to have shown just that. So a significant challenge to the orthodox view of language can only be laid to rest once we have an adequate and fully explanatory theory of metaphor in hand.

Second, even if we grant that the study of metaphor is not a distinctively philosophical matter, still metaphor is an interesting and pervasive feature of language use; and a distinctively philosophical investigation can still contribute significantly to understanding it. By thinking at a relatively abstract level about metaphorical utterances and the cognitive capacities necessary for comprehending them, and by bringing to bear the general constraints on language use that have been uncovered by 'core' analytic philosophy of language, a philosophical treatment of metaphor may succeed in developing fruitful hypotheses for empirical investigation where psychologists, cognitive scientists, and sociologists would not.

These two points suffice to demonstrate that philosophy of language has both a need and a right to investigate metaphor. But I also think we can, and should, say more. I think the study of metaphor demonstrates the importance of certain patterns and processes of thought which are almost entirely neglected by philosophy of mind. And I think this neglect significantly distorts philosophical investigations into the mind more generally. Specifically, as I have repeatedly reiterated, I think the study of metaphor demonstrates the unsustainability of drawing too sharp a distinction between purely rational and merely causal mental processes.

The patterns and processes of thought on which metaphor depends rely much more heavily on peculiarly human interests and capacities than do the paradigmatic cognitive representation, belief, and the paradigmatic rational relation, deductive inference. As we have seen, many philosophers have assumed as a consequence that these patterns and processes of thought must therefore be counted as merely causal, idiosyncratic 'associations'. But as I have argued, this assumption is incorrect. In order to do full justice to our intuitions about characterizations and aspectual thought, we need to acknowledge that these mental representations and processes are still genuinely rational, normative, and objective.

291

Further, once we understand how metaphor relies on patterns and processes of thought that exhibit a peculiarly human rationality, we can also see that much of the rest of our mental lives and uses of language work in this way as well.

Many philosophers, of course, think that rational law is a species of causal law, and so they accept that there is a sort of overlap between the rational and the causal. However, they still accept the crucial assumption that only some mental states and transitions are appropriately described — and evaluated — in rational terms, and that there is a clean distinction between those that are appropriately so described and those that are not. When mental representations and processes are governed by rational law, the thought goes, then a representation's content can be fully captured in propositional terms. A mental representation is factually correct when its propositional content matches the way the world is, and epistemically correct when it is formed in accord with an appropriate, rationally acceptable principle (perhaps the subject needs to be aware of this principle, perhaps not). A mental process, in turn, is correct when it occurs in accord with a rule that preserves propositional correctness. Philosophy is concerned to analyze the form and logic of such propositional representations and transitions, and their consequences for how we ought to think and act. By contrast, when mental representations and processes are governed by the merely causal process of association or reminding, then there may be important generalizations about which transitions between mental states a thinker in fact makes on different occasions, and which transitions different thinkers make on relevantly similar occasions. However, these generalizations are a subject for empirical investigation rather than for philosophical reflection.

I certainly think that these two sorts of law do govern major portions of our mental lives. On the one hand, we do reason about propositions that we believe or hypothesize to be true, we do intend for our reasoning to be valid, and we do feel compelled to correct any errors we make. On the other hand, we do let our minds wander from image to memory to thought in a merely associative way. In these cases, nothing like correctness enters into the picture. However, I also think that we often have mental states which deserve to be considered representational but which are not fully propositional — at least, not on the common, sentential understanding of that term. The commitment to a certain emotion being a fitting

response to someone, given what one takes his character to be, is a prime example of this. And I think we often engage in patterns of thought which deserve to be called rational but which cannot be fully subsumed within a formal logic, even a higher-order logic of relatively great expressive power that accepts induction as an inferential rule. The process of feature introduction through aspect application is a prime example of this.

These thoughts and transitions are considerably richer and messier than the forms of practical and theoretical reasoning on which philosophers usually focus. (Typical examples of the latter might be forming the intention to go to the fridge on the basis of my belief that there's beer in the fridge and my desire to drink a beer; or coming to the conclusion that there must still be beer in the fridge on the basis of my beliefs that there was beer there last Friday, when I looked, and that no one has been home in the intervening time.) As such, they may not necessarily be the best places to begin constructing a philosophical theory of the mind. But as exercises of our human rationality, these thoughts and transitions deserve to be treated within a complete philosophical theory of the mind. A theory that cannot treat them is for that very reason inadequate.

Much the same point goes for the philosophy of language. In the 1930's and 1940's, in Cambridge and especially Vienna, it seemed right to dismiss all of what we would now call pragmatics as lying outside the ambit of a distinctively rational cognitive capacity. Philosophy might hope to construct an ideal, formal language — a *Begriffschrift* or a logical calculus of the sort envisioned by Leibniz, Frege, Russell, and Carnap — into which we might translate the cognitive, propositional contents of thoughts and utterances, and thereby lay bare the rational relations obtaining among those contents. Philosophy might also hope to explain the conditions under which words refer to objects and properties. But the ends to which people actually put language seemed at that time to be too multifarious and too fully embedded in concrete situations involving particular people's particular interests, emotions, and beliefs to be susceptible of systematic, *a priori* description. Austin's work on speech acts in ordinary language successfully challenged this narrow view. And Grice's work on meaning in particular is so important, and has been so influential, because it laid out a framework for treating language *use* in rational,

293

systematic terms as well, and so for developing a distinctively philosophical account of pragmatics. It

thereby allowed us to expand our conception of philosophy of language considerably, so that few

philosophers would now be willing to accept such a heavily circumscribed view of philosophy.

In "Logic and Conversation," Grice at one point asks why we should think that language users

actually employ his Cooperative Principle and its attendant Maxims. He responds:

> A dull but, no doubt at a certain level, adequate answer is that it is just a well-recognized empirical fact
> that people behave in these ways; they learned to do so in childhood and have not lost the habit of doing
> so; and indeed it would involve a good deal of effort to make a radical departure from the habit.... I am,
> however, enough of a rationalist to want to find a basis that underlies these facts, undeniable though they
> may be; I would like to be able to think of the standard type of conversational practice not merely as
> something that all or most do *in fact* follow but as something that it is *reasonable* for us to follow, that we
> *should not* abandon.[32]

I, along with many other philosophers, think that Grice's theoretical framework has provided us with a

reason for accepting his view: that conversational practice in general is a rational practice, governed by

normative principles for coordinating our behavior so as to achieve common communicative ends.

Searle's treatment of linguistic communication as a species of rational action provides us with more

detailed analyses along similar theoretical lines. Notice here that Grice and Searle have continued to

accept the assumption that philosophy is essentially tied up with rational justification; they have simply

uncovered a way to extend the range of rational justification further than was once thought possible.

But it is important to remember here that the sorts of rational reconstructions that Griceans and

Searleans offer to explain why conversational practice is indeed rational, both in general and for particular

cases, never meet the high standard of a fully explicit, formalizable *calculus ratiocinatur* — even though

Grice himself sometimes makes it sound like they must.[33] They always appeal to 'the most likely

explanation', given the thinker's other assumptions about the conversational context and in particular her

assumptions about her interlocutor's interests and assumptions. And the ability to discern an

interlocutor's state of mind requires knowledge of human psychology, not just of the laws of pure reason.

This is true not just for explaining how we can understand metaphor and other indirect speech acts, but

---

[32] "Logic and Conversation," p. 29, emphasis in original.
[33] Cf. his discussion of the 'principle of calculability' for conversational implicature.

also for explaining how we understand direct, literal speech. Even here, the hearer must reason that the speaker most likely means just what she's saying, given the conventional meanings of the words uttered and the larger context of utterance.[34]

I suspect most philosophers of language would readily grant this point about the looseness of the proffered rational reconstructions, but insist that it does not call into question Grice's justification of the rational status of conversational practice. They would respond that this difficulty merely concerns the appropriate selection of *premises* about the interlocutor and the conversational context, and that this remains a merely psychological and sociological matter. Grice has performed an important task by showing us how a hearer can calculate the speaker's meaning *given* such premises, how a speaker can calculate that a hearer will calculate thus, and thus why a speaker has good reason to speak as she does. Lewis's and Schiffer's discussions of conventionality might be seen in the same way:[35] they analyze the rational assumptions and inferences that explain how conventions arise in the first place, but they don't tell us anything about which particular conventions will arise.

Other philosophers and theorists, such as Fodor or Chomsky, might accept also this point about the importance of human psychology to such rational reconstructions, but respond by returning to the earlier, pre-Gricean position. They might say, 'So much the worse for pragmatics — or even semantics. All we can treat are the rules of mental calculation themselves, and perhaps how they can be manifested in the brain. There are no real rules governing how these mental rules in fact apply to anything as messy as conversation.'[36]

I think both these responses concede too much. The task for the philosophy of language and of mind which results from these concessions is so impoverished as to be irrelevant. One important and distinctive task for philosophy of language is to develop a formal language which is as expressive as possible and which most closely approximates the surface syntax of natural languages while still retaining its formality. And one important and distinctive task for philosophy of mind is to explain how a brain

---

[34] Cf. Bach and Harnish, *Linguistic Communication and Speech Acts*, p. 63.
[35] Cf. David Lewis, *Convention: A Philosophical Study*; Stephen Schiffer, *Meaning*.
[36] Cf. Fodor, "Methodological Solipsism Considered as a Research Strategy in Cognitive Psychology."

which encodes the syntax of such a language could count as thinking and in particular as responsible for its thoughts and actions. But these are not the only tasks for philosophy. Another important job is attempting to make sense of — in the sense of rationally explain — the possibility and utility of our existing communicative practices. And if this task requires treading further into the territory of psychology than many philosophers are comfortable going, then so much the worse for the boundary between psychology and philosophy. Rationality is not always a matter of pure reason, and an explanation of the full range of our communicative practices requires appealing to the full range of human rationality.

In this dissertation, I have tried to do for metaphor specifically what Grice wanted to do for 'standard conversational practice' generally: not just to describe the regularities which *in fact* make metaphorical communication possible, but to show how metaphorical communication is "something that it is *reasonable* for us to follow, that we *should not* abandon." One important step along this path is establishing the basic claim above: showing that metaphorical utterances can indeed be treated in a broadly Gricean, pragmatic fashion.

It's also important to point out, however, that by examining metaphor closely, we discover that Grice was actually *wrong* to assume that metaphor could be treated as simply a species of implicature. Metaphor, like figurative language generally, differs from implicature in several significant ways. As I've noted, figurative speech acts always inherit the constraints on illocutionary force fixed by the grammatical mood of the sentence uttered; they can serve as the basis for further indirection; and it is natural for hearers to take up the speaker's figurative use of the words uttered in response. We could not have ascertained these differences between figurative language and implicature without a close examination of the phenomena themselves.

In addition to showing that metaphor can be accommodated within the broad framework of the pragmatic conversational model, and thus establishing that metaphorical utterances inherit the rational justification that Grice provided for conversation in general, I've also tried to push the rational justification one step further. In laying out what Searle calls the "specific shared system of principles" on

the basis of which we can "provide a rational reconstruction of the inference patterns that underlie our ability to understand" metaphors,[37] I have tried to show that this "shared system of principles" is *itself* in turn rational and objectively-oriented, and not just that it serves as the causal mechanism for a higher-order practice which is rational. Given this, the "system of principles" itself deserves a place within the philosophy of mind.

I claim that the relevant "system of principles" relies on aspectual thought: metaphorical communication exploits aspectual thought to restructure our characterization of the subject under consideration, and thereby to isolate certain features as those the speaker intends to predicate of the subject under discussion. I think three distinct elements within this system should be count as rational. I have argued, first, that characterizations are in the business of objective representation: that they are assessable for how accurately and appropriately they represent their subjects, in virtue of which constituents they include and in virtue of the structure in which they locate those constituents relative to one another. Second, I have argued that aspectual thought restructures those characterizations in a principled way. Although there are important stylistic and merely empirical, especially contextual, variations in the details of aspect-application, nonetheless thinkers can still have reasons for making certain matches between the two characterizations in advance of or instead of others. Third and finally, because it is often exceedingly difficult to make explicit the representational commitments which are embodied in our characterizations (in particular, the structural relations of prominence and centrality), and because metaphor so centrally exploits characterizations, metaphor provides an especially efficient tool for communicating these commitments. Thus, metaphorical speech plays an important, useful, and distinctive role in our communicative practices. It is a pattern of speech in which we have specific reason to engage, because it enables us to express representational commitments which we would otherwise find it very difficult to make explicit.

The points in these last four paragraphs — that our practice of speaking metaphorically inherits the rational justification which Grice provided for conversational practice in general, and that it receives a

---

[37] "Metaphor," p. 104.

specific justification from its distinctive role in communicating certain complex and nuanced, but objectively-oriented mental representations — further establish the distinctively philosophical relevance of a study of metaphor. But they should not be taken to show that metaphorical speech, or the patterns and processes of thought on which it depends, have just the same rational status as literal speech and propositional reasoning do. I think is also important to recognize that characterizations and aspectual thought are more fully embroiled in the particularities of a distinctively human psychology than standard propositional thought and deductive reasoning are. Three points in particular are relevant here.

First, concepts form the basic building blocks of thought, without which propositional representation might arguably not be possible at all. By contrast, characterizations are an ancillary, more sophisticated cognitive resource. A thinker who lacked characterizations could still be entirely competent as a rational agent in the strictest sense of that term. She could still keep track of objects and properties in the world, deduce further conclusions from her existing body of beliefs, and exchange information and engage in rational arguments with others.

Such a thinker would likely be seriously impaired when it came to the nuances of communicating and negotiating her way around the human world in general, not just in comprehending metaphors. If she lacked characterizations but still had concepts, this would presumably be because she operated with unstructured sets of beliefs and other attitudes about the objects and properties around her, rather than with structured constellations of those beliefs and attitudes. And because the structures of prominence and centrality play such important roles in our cognitive lives, lacking characterizations would make it more difficult to deal with other agents. In particular, a shared sense of the relative prominence of various objects, features, and facts greatly facilitates smooth communication and coordinated purposive activity more generally. And intuitions about the relative centrality of different features to the identity of individuals and kinds play an important role in predicting the behavior of those individuals and kinds through changes in which features they possess. Nonetheless, such a thinker would still have her distinctively conceptual capacities intact.

Second, if we turn from mental representations to mental processes, inference is the basic connecting glue of thought. By contrast, aspect-application is a much more specialized capacity, and one which calls for 'genius', as Davidson says. The facts that experimental subjects find it so difficult to perceive analogies even when they are contextually salient and instrumentally useful, and that experts in a given field find it much easier to perceive patterns and structural similarities, suggest that aspect-application more generally requires specific sensitivities and skills that not all thinkers possess. At the same time, the payoff that perceiving an analogy can bring for solving practical and theoretical problems suggests that aspect application more generally is a useful cognitive tool.

Third, if we turn from aspectual thought to metaphors specifically, the ability to generate and comprehend metaphors — especially richer and more novel metaphors — is not in principle an essential communicative skill. There is some empirical evidence that hearers with cognitive disorders that impair their ability to represent other people's mental states are also impaired in their comprehension of metaphors without thereby being (as severely) impaired in their comprehension of literal speech.[38] This suggests that literal comprehension is a more basic skill than metaphorical comprehension. Finally, *generating* metaphors, especially novel ones, requires more sophisticated and creative cognitive abilities than mere metaphorical comprehension does. One can be an insightful and sensitive reader or critic without oneself being a poet.[39] It does not seem that the capacities to make and to comprehend literal speech do or even could come apart in this way.

Given these points, I think we should draw two final lessons from the study of metaphor for philosophical investigations into the mind and language. The first point is that the patterns and processes of thought on which metaphor relies are not fundamental building blocks of thought, but are still important and genuinely rational human cognitive capacities. We are rational animals in a traditional,

---

[38] Cf. Robyn Langdon et al., "Understanding Minds and Understanding Communicated Meanings in Schizophrenia"; F. G. E. Happé, "Understanding Minds and Metaphors: Insights From the Study of Figurative Language in Autism."
[39] Cf. Arthur Danto, "Metaphor and Cognition."

more narrow sense of that phrase. We do make inferential and evidential transitions from one thought to another. Even when we don't make those transitions explicitly, we do hold ourselves and each other responsible for the truth of our claims and the validity of our reasoning. Our capacity to reason in this way enables us to build bridges that stay up, to govern ourselves in political and economic communities, and to make and justify ethical decisions. However, those are not the only sorts of mental activities in which we engage, and our capacity to reason according to a logical calculus is not segregated from our other mental capacities, even when we are engaged in these sorts of mental activities. We are rational animals in a richer sense of that phrase as well.

The second lesson is that the rest of language use is more like metaphorical communication than philosophers usually like to admit. This should not be surprising, given the fact that language use is designed to expresses and communicate our thoughts, and so reflects and relies upon the patterns of thought in which we typically engage. In particular, we have seen that many of the complaints (and praises) made about metaphor, which have been taken to demonstrate that metaphor is a deeply non-cognitive and non-propositional phenomenon, are in fact also exhibited by literal speech, to a greater or lesser degree. Metaphor is taken to express rough rather than determinate propositional content, to convey structured complexes of propositions rather than single claims, to be picture-like in conveying highly specific information, to reflect and to arouse affective attitudes, to rely upon and reinforce stereotypes. But each of these points, and often all of them together, can be true of literal speech as well. Much of our literal talk is rough, loose, complex, evocative and emotionally expressive.[40] We invoke characterizations and we intend to get our hearers to think of our subjects under new aspects, all while still speaking literally. The difference between literal and metaphorical speech is thus in significant part a matter of degree: most metaphors do tend to exhibit the cited features more palpably than most literal speech does. The only essential difference between metaphorical and literal speech, then, lies in the

---

[40] Some people have taken the fact that looseness, evocativeness, and so on are such pervasive features of language to demonstrate that these features are semantic rather than pragmatic. I think, on the contrary, that the very fact that we recognize that such speech is *merely* loose, evocative, etc., rather than explicit, precise, etc., demonstrates that they should be treated pragmatically.

This difference is that only in comprehending metaphors is it necessary to cultivate an aspect in order to determine the content of the speaker's *primary* communicative intention.

One great success of analytic philosophy has been the development and application of a rigorous, formal logic for representing thoughts and justified transitions between them. But this great success has also tended to stifle philosophical investigation into those areas of the mind and language that don't fit this model. Throughout the course of analytic philosophy, philosophers have tended to think of the mind under the aspect of a computing machine, and of language as a symbol system for interfacing between that machine and the world around it; some philosophers have even come to believe that this model applies in a fully literal way. In the terms I have set out in the last five chapters, this model for thinking about thought has brought both benefits and dangers. It has highlighted certain features of our behavior and our mental lives: our capacity for systematic reasoning over representations that are partially constituted by a complex syntax, and for communicating with others by exploiting our shared capacity to reason in this way. It has thereby helped us to recognize the centrality of concept possession and manipulation for explaining many of our other mental capacities.

But like any aspect, the aspect of the computing machine has also led us to neglect those features of the mind and of language use which don't find matches in the governing characterization. In particular, it has downplayed mental representations which do not take a sentential form — affective attitudes, images, readinesses for action, and characterizations, among others — and mental processes which are not deductive, but abductive or imaginative. To the extent that we philosophers do attend to these neglected features at all, we tend to treat them as belonging to separate, distinctively non-cognitive faculties. In reality, though, our cognitive, affective, imaginative, and kinesthetic faculties are all integrated into a single self. Our easy and fluid generation and appreciation of metaphors provides clear evidence of this, because it both depends upon and enables us to express the complex, structured representations that are produced by such a complex, integrated self.

BIBLIOGRAPHY

Armstrong, Sharon Lee, Lila R. Gleitman, and Henry Gleitman, "What Some Concepts May Not Be," *Cognition* 13, 1983; reprinted in *Concepts: Core Readings*, ed. Eric Margolis and Stephen Laurence (Cambridge, MA: MIT Press, 2000).

Austin, J. L.: *How to Do Things with Words*, 2nd edition, ed. J. O. Urmson and M. Sbisà (Cambridge, MA: Harvard University Press, 1962).

Bach, Kent: "Speaking Loosely," in *Figurative Language* (Midwest Studies in Philosophy XXV), ed. P. French and H. Wettstein (Oxford: Blackwell Publishers, 2001), pp. 249-263.

———: "Seemingly Semantic Intuitions," in *Meaning and Truth: Investigation in Philosophical Semantics*, ed. J. K.Campbell, M. O'Rourke, and D. Shier (New York: Seven Bridges Press, 2001).

———: "The Semantics/Pragmatics Distinction: What It Is and Why It Matters," in *The Semantics/Pragmatics Interface from Different Points of View*, ed. K. Turner (Oxford: Elsevier, 1999).

Bach, Kent and Robert Harnish: *Linguistic Communication and Speech Acts* (Cambridge, MA: MIT Press, 1979).

Barsalou, Lawrence: "Intraconcept Similarity and Its Implications for Interconcept Similarity," in *Similarity and Analogical Reasoning*, ed. A. Ortony and S. Vosniadou (Cambridge: Cambridge University Press, 1989).

Beardsley, Monroe: "The Metaphorical Twist," *Philosophy and Phenomenological Research* 22:3 (March 1962), reprinted in *The Aesthetic Point of View: Selected Essays*, ed. M. Wreen and D. Callen (Ithaca: Cornell University Press, 1982).

———: *Aesthetics: Problems in the Philosophy of Criticism* (New York: Harcourt, Brace and World, 1958).

Benjamin, Walter: "The Writer's Technique in Thirteen Theses," in *One-Way Street and Other Writings* (London: Verso Books, 1997).

Bergmann, Merrie: "Metaphorical Assertions," in *Pragmatics*, ed. S. Davis (Oxford: Oxford University Press, 1991).

Bezuidenhout, Anne: "Metaphor and What is Said: A Defense of a Direct Expression View of Metaphor," in *Figurative Language* (Midwest Studies in Philosophy XXV), ed. P. French and H. Wettstein (Oxford: Blackwell Publishers, 2001), pp. 156-186.

Black, Max: "How Metaphors Work: A Reply to Donald Davidson" in *On Metaphor*, ed. S. Sacks (Chicago: University of Chicago Press, 1978).

———: "Metaphor," in *Models and Metaphors: Studies in the Philosophy of Language* (Ithaca: Cornell University Press, 1962).

Blackburn, Simon: *Spreading the Word: Groundings in the Philosophy of Language* (Oxford: Oxford University Press, 1984).

Booth, Wayne: "Metaphor as Rhetoric: The Problem of Evaluation," in *On Metaphor*, ed. S. Sacks (Chicago: The University of Chicago Press, 1978).

Borg, Emma: "An Expedition Abroad: Metaphor, Thought, and Reporting", in *Figurative Language* (Midwest Studies in Philosophy XXV), ed. P. French and H. Wettstein (Oxford: Blackwell Publishers, 2001), pp. 227-248.

Boyd, Richard: "Metaphor and Theory Change: What is 'Metaphor' a Metaphor For?" in *Metaphor and Thought*, ed. A. Ortony (Cambridge: Cambridge University Press, 1979).

Brandom, Robert: *Making it Explicit: Reasoning, Representing, and Discursive Commitment* (Cambridge, MA: Harvard University Press, 1994).

Brooks, Cleanth: "The Heresy of Paraphrase," in *The Well Wrought Urn* (New York: Harcourt Brace, 1947).

Camp, Elisabeth: "The Generality Constraint, Nonsense, and Categorial Restrictions," *Philosophical Quarterly*, forthcoming.

———: "The Conceptual but Non-Propositional Content of Perception," manuscript.

Cappelen, Herman and Ernie LePore, "On an Alleged Connection between Indirect Speech and the Theory of Meaning," *Mind and Language* 12:3/4 (September/December 1997), pp. 278-296.

Carston, Robyn: "The Semantics/Pragmatics Distinction: A View from Relevance Theory," in *The Semantics/Pragmatics Interface from Different Points of View*, ed. K. Turner (Oxford: Elsevier, 1999).

Cavell, Stanley: "Aesthetic Problems of Modern Philosophy," in *Must We Mean What We Say?* (Cambridge: Cambridge University Press, 1976).

Churchland, Paul: *Scientific Realism and the Plasticity of Mind* (Cambridge: Cambridge University Press, 1979).

Cohen, L. Jonathan and Avashai Margalit: "The Role of Inductive Reasoning in the Interpretation of Metaphor," in *Semantics of Natural Language*, ed. D. Davidson and G. Harman (Dordrecht: D. Reidel, 1972).

Cohen, Ted: "Metaphor, Feeling, and Narrative," *Philosophy and Literature* 21 (1997), pp. 223-244.

———: "Metaphor and the Cultivation of Intimacy," in *On Metaphor*, ed. Sheldon Sacks (Chicago: The University of Chicago Press, 1978).

———: "Notes on Metaphor," *Journal of Aesthetics and Art Criticsm* 34 (Spring 1976), pp. 249-259.

———: "Figurative Speech and Figurative Acts," *Journal of Philosophy* 71 (1975), pp. 669-684.

Conrad, Joseph: *Heart of Darkness*, ed. J. Woolley (New York: Penguin USA, 1999).

Crimmins, Mark: "Hesperus and Phosphorus: Sense, Pretense, and Reference," *Philosophical Review* 107:1 (January 1998), pp. 1-47.

Dagut, M.B.: "Can 'Metaphor' be Translated?" *Babel* 22 (1976), pp. 21-33.

Danto, Arthur: "Metaphor and Cognition," in *Metaphor and Knowledge: Knowledge and Language*, vol. III, ed. F.R. Ankersmit and J.J.A. Mooij (Dordrecht: Kluwer Academic Publishers, 1993).

Davidson, Donald: "What Metaphors Mean," in *Inquiries into Truth and Interpretation* (Oxford: Clarendon Press, 1984).

———: "Nice Derangement of Epitaphs", in *Truth and Interpretation: Perspectives on the Philosophy of Donald Davidson*, ed. E. LePore (Oxford: Blackwell, 1986).

———: "A Coherence Theory of Truth and Knowledge," in *Truth and Interpretation: Perspectives on the Philosophy of Donald Davidson*, ed. E. LePore (Oxford: Blackwell, 1986).

Davies, Martin: "Idiom and Metaphor," *Proceedings of the Aristotelian Society* 83 (1983), pp. 67-85.

Dennett, Daniel: "True Believers: The Intentional Strategy and Why It Works," in *Mind Design II: Philosophy, Psychology, Artificial Intelligence*, ed. J. Haugeland (Cambridge: MIT Press, 1997).

———: "Brain Writing and Mind Reading," in *Brainstorms: Philosophical Essays on Mind and Psychology* (Montgomery VT: Bradford Books, 1978).

DeRose, Keith: "Assertion, Knowledge and Context," *Philosphical Review,* forthcoming.

Dostoyevsky, Fyodor: *Crime and Punishment* (New York: Modern Library, 1950).

Dretske, Fred: "The Nature of Thought," in *Perception, Knowledge and Belief: Selected Essays* (Cambridge: Cambridge University Press, 2000).

Dreyfus, Hubert: "From Socrates to Expert Systems: The Limits and Dangers of Calculative Rationality," manuscript, available at: http://ist-socrates.berkeley.edu/~hdreyfus/html/papers.html.

Eliot, T.S.: "Love Song of J. Alfred Prufrock," in *Collected Poems, 1909-1962* (London: Faber and Faber, 1974).

Empson, William: *The Structure of Complex Words* (Norfolk, CN: New Directions, 1951).

Evans, Gareth: *The Varieties of Reference,* ed. J. McDowell (Oxford: Clarendon Press, 1982).

Fodor, Jerry: "Methodological Solipsism Considered as a Research Strategy in Cognitive Psychology," *Behavioral and Brain Sciences* 3 (1980), pp. 63-109.

Fodor, Jerry and Ernie LePore: "The Red Herring and the Pet Fish: Why Concepts Still Can't be Prototypes," *Cognition* 58:2 (1996), pp 253-270.

Fogelin, Robert: *Figuratively Speaking* (New Haven: Yale University Press, 1988).

Frege, Gottlob: "Thought," in *The Frege Reader* ed. M. Beaney (Oxford: Blackwell, 1997).

Gass, William: "In Terms of the Toenail: Fiction and the Figures of Life," in *Fiction and the Figures of Life* (Boston: Nonpareil Books, 1979).

Gentner, Dedre: "The Mechanisms of Analogical Learning," in *Similarity and Analogical Reasoning*, ed. A. Ortony and S. Vosniadou (Cambridge: Cambridge University Press, 1989).

Gibbard, Allan: *Wise Choices, Apt Feelings: A Theory of Normative Judgment* (Cambridge: Harvard University Press, 1990).

Gibbs, Raymond: *The Poetics of Mind* (Cambridge: Cambridge University Press, 1994).

Goatly, Andrew: *The Language of Metaphors:* (London and New York: Routledge, 1997).

Goodman, Nelson: *Languages of Art: An Approach to a Theory of Symbols* (New York: Bobbs-Merrill Company, 1968).

Grice, Paul: "Logic and Conversation," in *Studies in the Way of Words* (Cambridge, MA: Harvard University Press, 1989).

Happé, F. G. E.: "Understanding Minds and Metaphors: Insights from the Study of Figurative Language in Autism," *Metaphor and Symbolic Activity* 10 (1995), pp. 275-95.

Harman, Gil: "(Nonsolipsistic) Conceptual Role Semantics," *Reasoning, Meaning, and Mind* (Oxford: Clarendon Press, 1999).

Hauser, Marc: *Wild Minds: What Animals Really Think* (New York: Henry Holt and Co., 2000).

Hausman, Carl: "Metaphors, Referents, and Individuality," *Journal of Aesthetics and Art Criticism* 42:2 (Winter 1983), pp. 181-196.

Henle, Paul: "Metaphor," in *Language, Thought, and Culture* (Ann Arbor: University of Michigan Press, 1958), pp. 173-195.

Hills, David: "Aptness and Truth in Verbal Metaphor," *Philosophical Topics* 25:1 (Spring 1997), pp. 117-153.

Hobbes, Thomas: *Leviathan,* ed. A.P. Martinich (Peterborough, Ontario: Broadview Press, 2002).

Holyoak, Keith: "The Pragmatics of Analogical Transfer," in *The Psychology of Learning and Motivation*, vol. 19, ed. G. H. Bower (New York: Academic Press, 1985).

Holyoak, Keith, and Paul Thagard: *Mental Leaps: Analogy in Creative Thought* (Cambridge: MIT Press, 1995).

Homer: *The Iliad*, trans. R. Fitzgerald (Oxford: Oxford University Press, 1998).

Hopkins, Robert: *Picture, Image and Experience: A Philosophical Inquiry* (Cambridge: Cambridge University Press, 1998).

Hornsby, Jennifer: "Meaning and Uselessness: How to Think about Derogatory Words," in *Figurative Language* (Midwest Studies in Philosophy XXV), ed. P. French and H. Wettstein (Oxford: Blackwell Publishers, 2001), pp. 128-141.

Johnson-Laird, Philip: "Analogy and the Exercise of Creativity," in *Similarity and Analogical Reasoning*, ed. A. Ortony and S. Vosniadou (Cambridge: Cambridge University Press, 1989).

Kant, Immanuel: *Critique of Judgement*, trans. J. C. Meredith (Oxford: Clarendon Press, 1952).

Kaplan, David: "Quantifying In," in *Reference and Modality*, ed. L. Linsky. (Oxford: Oxford University Press, 1971).

———: "Dthat," in *Pragmatics* (*Syntax and Semantics*, no. 9), ed. P. Cole (New York: Academic Press, 1978).

———: "Afterthoughts," in *Themes from Kaplan*, ed. J. Almog, J. Perry, and H. Wettstein (Oxford: Oxford University Press, 1989).

———: "The Meaning of 'Ouch' and 'Oops'," manuscript (presented at Arché Research Centre, at the University of St. Andrews; and at other locales.)

Katz, Jerrold: *The Philosophy of Language* (New York: Harper and Row, 1966).

Katz, Jerrold and Fodor, Jerry: "The Structure of a Semantic Theory," in *The Structure of Language*, ed. J. Fodor and J. Katz (Englewood Cliffs, NJ: Prentice Hall, 1964).

Kelly, Sean: "The Non-Conceptual Content of Perceptual Experience: Situation Dependence and Fineness of Grain," *Philosophy and Phenomenological Research* LXII:3 (May 2001), pp. 601-608.

Kemp, G. Neville: "Metaphor and Aspect-Perception," *Analysis* 51:2 (March 1991), pp. 84-90.

King, Jeffery C.: "Structured Propositions and Complex Predicates," *Nous* 29:4 (1995), pp. 516-535.

King, Jeffrey C. and Jason Stanley, "Semantics, Pragmatics, and the Role of Semantic Content," forthcoming in *Semantics vs. Pragmatics*, ed. Z. Szabo (Oxford: Oxford University Press).

Kittay, Eva F.: *Metaphor: Its Cognitive Force and Linguistic Structure* (Oxford: Clarendon Press, 1987).

Kripke, Saul: *Naming and Necessity* (Cambridge, MA: Harvard University Press, 1980).

Kuhn, Thomas: "Metaphor in Science," in *Metaphor and Thought*, ed. A. Ortony (Cambridge: Cambridge University Press, 1979).

Kupperman, Joel: *Character* (Oxford: Oxford University Press, 1991).

Lakoff, George: "Hedges: a Study in Meaning Criteria and the Logic of Fuzzy Concepts," *Journal of Philosophical Logic* 2 (1973), pp. 458-508.

Lakoff, George, and Mark Johnson: *Metaphors We Live By* (Chicago: University of Chicago Press, 1980).

———: *More than Cool Reason: A Field Guide to Poetic Metaphor* (Chicago: University of Chicago Press, 1989).

Langacker, Ronald W: *Foundations of Cognitive Grammar* (Stanford: Stanford University Press, 1987).

Langdon, Robyn, Martin Davies and Max Coltheart, "Understanding Minds and Understanding Communicated Meanings in Schizophrenia," *Mind and Language* 17:1/2 (February/April 2002) pp. 68-104.

Laurence, Stephen and Eric Margolis, "Introduction" to *Concepts: Core Readings*, ed. E. Margolis and S. Laurence (Cambridge, MA: MIT Press, 2000).

Levin, Samuel: *The Semantics of Metaphor* (Baltimore: Johns Hopkins University Press, 1977).

———: "Standard Approaches to Metaphor and a Proposal for Literary Metaphor," in *Metaphor and Thought*, ed. A. Ortony (Cambridge: Cambridge University Press, 1979).

———: *Metaphoric Worlds: Conceptions of a Romantic Nature* (New Haven: Yale University Press, 1988.)

Lewis, C. S.: *The Discarded Image: An Introduction to Medieval and Renaissance Literature* (Cambridge: Cambridge University Press (Canto), 1994).

Lewis, David: *Convention: A Philosophical Study* (Cambridge MA: Harvard University Press, 1969).

Locke, John: *An Essay Concerning Human Understanding*, ed. P. Nidditch (Oxford: Clarendon Press, 1989).

McDowell, John: "Aesthetic Value, Objectivity, and the Fabric of the World," in *Mind, Value, and Reality* (Cambridge, MA: Harvard University Press, 1998).

———: "Values and Secondary Qualities," in *Mind, Value, and Reality* (Cambridge, MA: Harvard University Press, 1998).

———: "Virtue and Reason," in *Mind, Value, and Reality* (Cambridge, MA: Harvard University Press, 1998).

———: *Mind and World* (Cambridge, MA: Harvard University Press, 1994).

Medin, Douglas, and Andrew Ortony: "Psychological essentialism," in *Similarity and Analogical Reasoning*, ed. A. Ortony and S. Vosniadou (Cambridge: Cambridge University Press, 1989).

Merleau-Ponty, Maurice: *Phenomenology of Perception*, trans. C. Smith (London: Routledge and Kegan Paul, 1962).

Millikan, Ruth Garrett: "Thoughts without Laws," in *White Queen Psychology and Other Essays for Alice* (Cambridge, MA: Bradford/MIT Press, 1993).

Minsky, Marvin, "A Framework for Representing Knowledge," in *The Psychology of Computer Vision*, ed. P.H. Winston (New York: McGraw-Hill, 1975).

Moran, Richard, "Seeing and Believing: Metaphor, Image, and Force," *Critical Inquiry* 16 (Autumn 1989), pp. 87-112.

Nehamas, Alexander: "An Essay on Beauty and Judgment," *Threepenny Review* (Winter 2000).

Nogales, Patti: *Metaphorically Speaking* (Stanford: CSLI Publications (Center for the Study of Language and Information), 1999).

Novitz, David: *Knowledge, Fiction, and Imagination* (Philadelphia: Temple University Press, 1987).

Nussbaum, Martha: *Love's Knowledge: Essays on Philosophy and Literature* (Oxford: Oxford University Press, 1990).

Ortony, Andrew: "Beyond Literal Similarity," *Psychological Review* 87 (1979), pp. 161-80.

Peacocke, Christopher: "Does Perception Have a Non-Conceptual Content?" *Journal of Philosophy* 98:5 (May 2001), pp. 239-264.

———: "Scenarios, Concepts, and Perception," in *The Contents of Experience*, ed. Tim Crane (Cambridge: Cambridge University Press, 1992).

———: *A Study of Concepts* (Cambridge, MA: MIT Press, 1992).

Quine, Willard, V. O.: "Two Dogmas of Empiricism," in *From a Logical Point of View: Nine Logico-Philosophical Essays* (Cambridge, MA: Harvard University Press, 1953).

Ransom, John Crowe: "Poetry: I. The Formal Analysis," *Kenyon Review* 9 (Summer 1947), pp. 436-56.

Recanati, François: "Literal/Nonliteral," in *Figurative Language* (Midwest Studies in Philosophy XXV), ed. P. French and H. Wettstein (Oxford: Blackwell Publishers, 2001), pp. 264-274.

Reddy, M. J.: "A Semantic Approach to Metaphor," in *Papers from the Fifth Regional Meeting of the Chicago Linguistics Society*, ed. R. I. Binnick et al (Department of Linguistics, University of Chicago, 1969).

Reimer, Marga: "Davidson on Metaphor," in *Figurative Language* (Midwest Studies in Philosophy XXV), ed. P. French and H. Wettstein (Oxford: Blackwell Publishers, 2001), pp. 142-155.

Richards, I.A. *The Philosophy of Rhetoric* (Oxford: Oxford University Press, 1936).

Rorty, Richard: "Unfamiliar Noises I: Hesse and Davidson on Metaphor," *Proceedings of the Aristotelian Society* (1987, supp. volume 61), pp. 283-296.

Rosch, Eleanor: "Principles of Categorization", in *Concepts: Core Readings*, ed. E. Margolis and S. Laurence (Cambridge MA: MIT Press, 1999).

Rumelhart, David E.: "Understanding and Summarizing Brief Stories," in *Basic Process in Reading: Perception and Comprehension*, ed. D. Laberge and S. J. Samuels (Hillsdale, N. J.: Lawrence Erlbaum, 1977), pp. 265-303.

Stephen Schiffer, *Meaning* (Oxford: Clarendon Press, 1972).

Scruton, Roger: *Art and Imagination: A Study in the Philosophy of Mind* (London: Metheun & Co., 1982).

Searle, John: "Intentionality, Consciousness and the Background," in *The Rediscovery of the Mind* (Cambridge, MA: Bradford/MIT Press, 1992).

———: *Intentionality: An Essay in the Philosophy of Mind* (Cambridge: Cambridge University Press, 1983).

———: "Metaphor," in *Expression and Meaning: Studies in the Theory of Speech Acts* (Cambridge: Cambridge University Press, 1979).

———: "Indirect Speech Acts," in *Expression and Meaning: Studies in the Theory of Speech Acts* (Cambridge: Cambridge University Press, 1979).

———: "Literal Meaning," in *Expression and Meaning: Studies in the Theory of Speech Acts* (Cambridge: Cambridge University Press, 1979).

———: *Speech Acts: An Essay in the Philosophy of Language* (Cambridge: Cambridge University Press, 1969).

Sellars, Wilfred: *Empiricism and the Philosophy of Mind* (Cambridge, MA: Harvard University Press, 1997).

Shakespeare, William: *The Riverside Shakespeare*, 2nd edition, ed. J. J. M. Tobin et al. (New York: Houghton Mifflin Co., 1997).

Shank, Roger and Robert Abelson, *Scripts, Plans, Goals and Understanding* (Hillsdale, NJ: Erlbaum, 1977).

Smith, Edward and Douglas Medin, *Categories and Concepts* (Cambridge, MA: Harvard University Press, 1981).

Soames, Scott: *Beyond Rigidity: The Unfinished Semantic Agenda of* Naming and Necessity (Oxford: Oxford University Press, 2002).

Sontag, Susan: *Illness as Metaphor and AIDS and Its Metaphors*, combined ed. (New York: Doubleday, 1990).

Sperber, Dan and Wilson, Deirdre: "Pragmatics, Modularity, and Mind-Reading," *Mind and Language* 17:1/2 (February/April 2002), pp. 3-23.

———: *Relevance: Communication and Cognition,* 2nd edition (Oxford: Blackwell, 1995).

———: "Rhetoric and Relevance," in *The Ends of Rhetoric: History, Theory, Practice,* ed. D. Wellbery and J. Bender (Stanford: Stanford University Press, 1990).

———: "Loose Talk," *Proceedings of the Aristotelian Society* 86 (1986), pp. 153-171.

———: "Irony and the Use-Mention Distinction," in *Radical Pragmatics*, ed. P. Cole (New York: Academic Press, 1981).

Stalnaker, Robert: "Assertion," in *Context and Content: Essays on Intentionality in Speech and Thought* (Oxford: Oxford University Press 1999).

———: "Pragmatics," in *Context and Content: Essays on Intentionality in Speech and Thought* (Oxford: Oxford University Press 1999).

———: "Pragmatic Presuppositions," in *Context and Content: Essays on Intentionality in Speech and Thought* (Oxford: Oxford University Press 1999).

Stern, Josef: *Metaphor in Context* (Cambridge, MA: MIT Press, 2000).

Strawson, P. F.: "Imagination and Perception," in *Freedom and Resentment and Other Essays* (London: Metheun and Co., 1974).

Tirrell, Lynne: "Seeing Metaphor as Seeing-as: Remarks on Davidson's Positive View of Metaphor," *Philosophical Investigations* 14:2 (April 1991), pp. 143-154.

Tirrell, Lynne: "Extending: the Structure of Metaphor," *Nous* 23 (March 1989), pp. 17-34.

Travis, Charles, "On Constraints of Generality," *Proceedings of the Aristotelian Society* (1994), pp. 165-188.

———: *Unshadowed Thought: Representation in Thought and Language* (Cambridge, MA: Harvard University Press, 2000).

Tversky, Amos: "Features of Similarity," *Psychological Review* 84:4 (1977), pp. 327-352.

Wallace, R. Jay: "Addiction as a Defect of the Will: Some Philosophical Reflections," *Law and Philosophy* 18 (1999), pp. 621-654.

Walton, Kendall: "Metaphor and Prop-oriented Make-Believe," *European Journal of Philosophy* 1 (1993), pp. 39-57.

———: *Mimesis as Make-Believe: On the Foundations of the Representational Arts* (Cambridge, MA: Harvard University Press, 1990).

Weinreich, Uriel: "Explorations in Semantic Theory," in *Current Trends in Linguistics;* vol. III, *Theoretical Foundations,* ed. T. A. Sebeok (The Hague: Mouton Press, 1966).

White, Roger: *The Structure of Metaphor: The Way the Language of Metaphor Works* (Oxford: Blackwell, 1996).

Wiggins, David: *Sameness and Substance* (Oxford: Blackwell, 1980).

Wittengstein, Ludwig: *Tractatus Logico-Philosophicus,* trans. D. F. Pears and B. F. McGuinness (London: Routledge & Kegan Paul, 1961).

———: *Philosophical Investigations,* trans. G. E. M. Anscombe (Oxford: Basil Blackwell, 1972).

Wollheim, Richard: "Seeing-as, Seeing-in, and Pictorial Representation," in *Art and Its Objects*, 2nd edition (Cambridge: Cambridge University Press (Canto), 1980).

———: "Correspondence, Projective Properties, and Expression," in *The Mind and its Depths* (Cambridge, MA: Harvard University Press, 1993).

———: *Art and Its Objects*, 2nd edition (Cambridge: Cambridge University Press (Canto), 1980).

———: *On the Emotions* (New Haven: Yale University Press, 1999).

Woodfield, Andrew: "Conceptions," *Mind* C:4 (October 1991), pp. 547-72.

Wordsworth, William: "I Wandered Lonely as a Cloud," in *English Romantic Writers*, ed. D. Perkins (New York: Harcourt Brace Jovanovich, 1967).

Yablo, Stephen: "Is Conceivability a Guide to Possibility?," *Philosophy and Phenomenological Research* 53:1 (March 1993), pp. 1-42.

Yeats, William Butler: "Byzantium," in *The Collected Poems of W. B. Yeats*, ed. R. Finneran (New York: Macmillan, 1989).

Zwicky, A. and J. Sadock: "Ambiguity Tests and How to Fail Them," in *Syntax and Semantics*, vol. 4, ed. J. Kimball (New York: Academic Press, 1975).