

Metaphor and That Certain ‘Je ne sais quoi’*
(to appear in Philosophical Studies)

Philosophers have traditionally inclined toward one of two opposite extremes when it comes to metaphor. On the one hand, partisans of metaphor have tended to believe that metaphors do something different in kind from literal utterances; it is a ‘heresy’, they think, either to deny that what metaphors do is genuinely cognitive, or to assume that it can be translated into literal terms. On the other hand, analytic philosophers have typically denied just this: they tend to assume that *if* metaphors express any genuine content at all, then that content can in principle be paraphrased into literal terms. They often conclude on this basis that metaphor is theoretically dispensable, and so that it poses no special challenges and affords no distinctive insights for the philosophy of language and mind.

In this paper, I want to steer between these two extremes. Metaphors don’t do anything different in kind from what can be done with literal speech. But this does not render metaphor theoretically dispensable or irrelevant. In certain circumstances, I will argue, metaphors can enable speakers to communicate contents that cannot be stated in fully literal and explicit terms. These cases thus serve as counterexamples to the “Principle of Expressibility,” the idea that whatever can be meant can be said. Indeed, I will argue, the point goes for cognition as well as communication: metaphors can sometimes provide us with our only cognitive access to certain properties. In the end, I think, thinking about metaphor is useful because it draws our attention to patterns and processes of thought that play a pervasive role in our ordinary thought and talk, and that can genuinely extend our basic communicative and cognitive resources.

1. What Is an Adequate Paraphrase?

* Elisabeth Camp, Society of Fellows, Harvard University; ecamp@fas.harvard.edu. Thanks to John Hawthorne, Peter Koellner, John MacFarlane, Daniel Nolan, Chris Pincock, Francois Recanati, John Searle, and Dmitri Tymoczko; thanks also to the Harvard M&E workshop, the Corridor Group, and participants at colloquia at the University of Michigan and at the University of California, Santa Cruz. Special thanks for Jim Pryor for the title.

Before turning to metaphor specifically, we need to settle what, in general, an adequate paraphrase would be. Everyone can agree that a paraphrase “says the same thing another way”; the difficulty is that both ordinary and theoretical usage underwrite multiple interpretations of this phrase. I believe, first, that that an adequate paraphrase should capture the content of the speaker’s intended illocutionary act: it should state how the world would have to be for the speaker’s claim (promise, etc.) to count as true (or otherwise satisfied). Speakers can claim multiple propositions by their utterances; when they do, all those propositions should be captured by the paraphrase, not just the proposition (if any) expressed by the sentence which was actually uttered. The paraphrase should not, however, include contents the speaker merely insinuated, or merely caused her hearer to entertain.

Second, I think an adequate paraphrase must state that content in a literal and explicit fashion. This means, I think, that the paraphrase should consist of a sentence whose semantic content is the same as the content of the speaker’s intended speech act; it should enable an otherwise linguistically competent speaker to understand the original utterance’s content simply in virtue of understanding the meanings of the paraphrasing sentence’s constituent terms and their mode of combination. I am understanding linguistic competence here to include a basic ability to assign values to semantically context-sensitive terms even when this requires appeal to the speaker’s intentions, but not to include the rich interpretative abilities we bring to bear in recovering implicatures and related pragmatic phenomena.¹

2. What Do Metaphors Do?

In speaking metaphorically, we often make claims and other speech acts with quite complex contents. So, for instance, by saying of your friend’s new boyfriend

(1) Steve is a sheepdog

¹ These abilities are partially (though not, I think entirely) overlapping. Only in the former case, though, does the semantics itself call for exercising the relevant ability. Thus, we can imagine linguistically competent interpreters who are constitutionally disposed to cease interpreting once they have completed all semantic interpretation. The question of how explicit a paraphrase must be to count as adequate can be framed, conversely, as how explicit the intended content must be for the hearer’s failure to recover it to count as impugning his linguistic competence. If the failure to recover some bit of content would not impugn the hearer’s linguistic competence, then a paraphrase should make that content explicit.

(Nogales 1999, 1), you might claim, among other things, that Steve is strong, loyal, calm, attentive, and adoring, but not very smart or interesting. Similarly, in uttering

(2) Juliet is the sun,

I think that Romeo claims, among other things, that Juliet is worthy of worship; that she is the focus of his thoughts and dreams; that she is an exemplar of goodness and beauty; that her goodness and beauty are natural and original, and far exceed those of the other ladies of Verona; and that her goodness nurtures Romeo and helps him to grow emotionally.

Because metaphorical utterances like (1) and (2) express such complex contents in so few words, they are highly efficient vehicles for communication. In communicating, we confront an “articulatory bottleneck”: the process of vocal articulation runs about four times slower than the surrounding mental processes (Levinson 2000; 6, 28). We therefore need to pack as much information as we can into each of the syllables we utter, and rely on our hearers to decode the full richness of our intended meaning from those syllables. Metaphors offer one convenient way to do this.

They also enable an important sort of *cognitive* efficiency, one that goes beyond just circumventing the limitations of our vocal cords. As Glucksberg and Keysar (1993, 421) say, metaphors present “a patterned complex of properties in one chunk.” I believe that this is because metaphorical communication exploits a general fact about our engagement with the world: we typically experience multiple properties instantiated together, first in particular individuals and again across individuals of certain kinds. We intuitively associate those properties in our thinking, forming what I call “characterizations” of individuals and kinds, which are roughly similar to what others have called “stereotypes” or “prototypes.” In any particular context in which a characterization is evoked, just a subset of the properties within it will be salient. By mentioning a certain individual or kind — sheepdogs, say, or the sun, or Napoleon — we can invoke the subset of its salient properties; and we can do so, crucially, without either the speaker or the hearer needing to identify each constituent element in the set individually. A metaphor can thereby exploit the whole characterization, bundled together into a single intuitive ‘chunk’, to predicate a complex content of the subject under discussion.

Because the ‘chunks’ communicated by metaphor are so intuitive, there’s a temptation to think they are “undifferentiated,” as Josef Stern (2000, 191) says. I think this is wrong, and importantly so. Metaphors’ contents are often highly complex, not just in the sense of having many components, but in having a rich structure among those components. Consider, for instance, Romeo’s utterance of (2) in the context in which he utters it.² I think it’s part of Romeo’s claim here that some of the features he ascribes to Juliet are more important than others. Her being gloriously beautiful is more important, for instance, than her helping him to grow emotionally. Thus, if you wanted to give someone a succinct summary of Romeo’s claim, you’d better cite the first property but might well leave out the second; by contrast, anyone who insisted on the latter at the expense of the former would have misinterpreted Romeo. It’s also part of Romeo’s claim, I think, that certain of the features he claims Juliet to possess are responsible for, or entail, others. For instance, Juliet is worthy of worship *because* she’s the prime exemplar of beauty and goodness. And because she’s the exemplar of these virtues, her beauty and goodness must themselves be original and genuine; the virtues possessed by the other ladies of Verona must, by contrast, be merely artificial and derivative. These claims form an interconnected network of implications, with Juliet’s being the exemplar of beauty and goodness at the center. Thus, if you wanted to give someone an interpretive key for unlocking how Romeo is claiming Juliet to be, you’d do better to cite her being an exemplar of beauty and goodness than to cite her being the focus of Romeo’s thoughts and dreams.

Metaphors are efficient vehicles for communicating such structured contents, I believe, because they crucially exploit characterizations. In forming a characterization of someone or something, we don’t just assemble a set of properties we believe the individual or kind characterized to possess. In general, some of those properties are intuitively more *prominent* in our thoughts than others, in the sense that they are more obvious and important, at least in that context. And I think we intuitively take some properties to be more *central* than others, in the sense that they cause or otherwise determine some of the other properties possessed by that individual or kind. Metaphors allow us to use the intuitive structure of our

² See the Appendix for the main text of Romeo’s speech.

characterization of one individual or kind to structure our thinking about something else, without requiring us to articulate or even to think about that structure itself explicitly.

3. Arguments Against “The Heresy of Paraphrase”

These observations about what metaphors do form the basis of the most well-established arguments against the idea that their contents can be given a literal translation. I’ll focus on just one representative articulation, by Max Black (1962, 46):

Suppose we try to state the cognitive content of a...metaphor in ‘plain language’. Up to a point, we may succeed...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; the relevant weakness of the literal paraphrase is not that it may be tiresomely prolix or boringly explicit (or deficient in qualities of style); it fails to be a translation because it fails to give the insight that the metaphor did.

Black makes at least three distinct, though related, complaints against “the heresy of paraphrase” here.

The first is that a paraphrase fails to capture the metaphor’s content, because it always says what it does “with the wrong emphasis,” by ‘leveling out’ the structure of what we understand when we understand a metaphor.³ I argued for the positive component of this claim in §2: metaphors’ contents often include a rich structure. However, I think Black’s negative claim is mistaken: a literal paraphrase need not ignore these structural relations. Our language may not contain many common, convenient devices for making them explicit, but it is well within our capabilities to represent them. I did just that with Romeo’s metaphor, albeit in rather laborious terms, in order to motivate the intuition that there *was* such a structure. We can also, more perspicuously, supplement our language with formal representational systems such as numerical weightings or maps, as I’ve done in Figure 1. (Here, the heaviness of the boxes indicate relative importance; the arrows indicate causes and implications among properties possessed.) Thus, this argument at most shows that the standard for an adequate paraphrase is higher than some have thought; it does not demonstrate that this standard is unattainable.

³ Cleanth Brooks (1947, 182) makes the same objection: that the paraphrase fails to capture “the ‘inner’ structure or the ‘essential’ structure or the ‘real’ structure” at least of poetic metaphors.

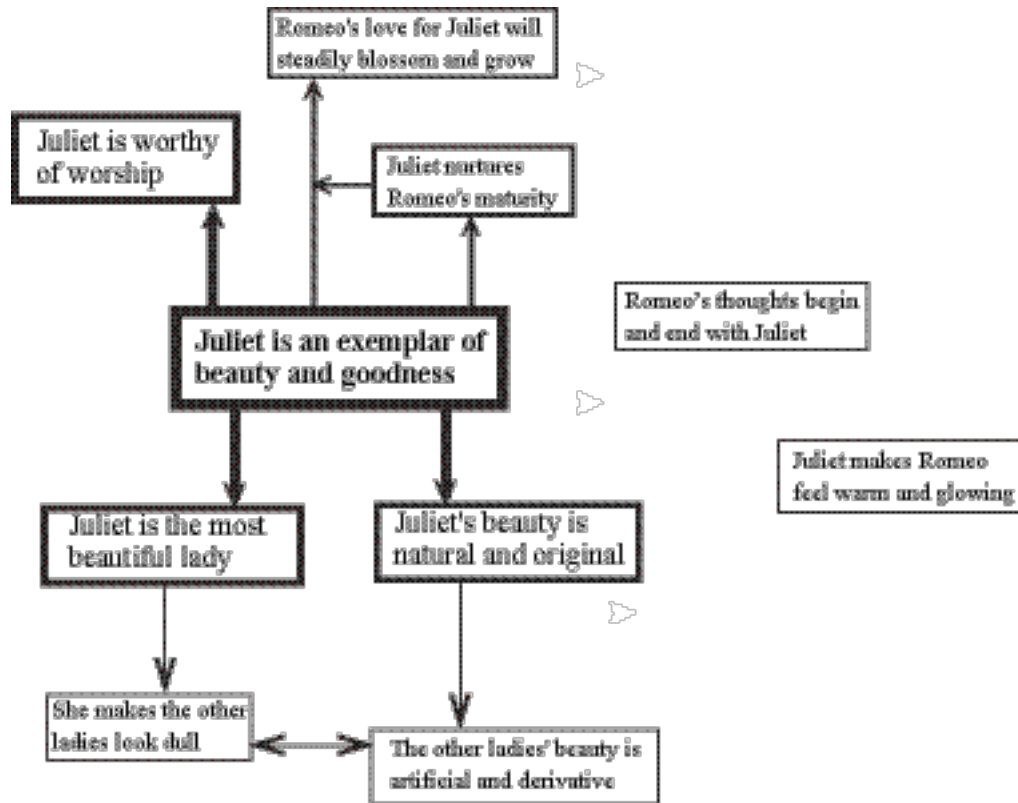


Figure 1.

Confronted with something like Figure 1, though, it's natural to respond with Black's second objection. Such a paraphrase fails to translate the metaphor because it "says too much": it is inappropriately specific, definite, and explicit. For instance, while Romeo plausibly meant something *like* that his thoughts begin and end with Juliet, it would be just as plausible to attribute alternative, related but distinct claims to him instead. And while the metaphor's content plausibly includes *some* structure of relative importance and centrality, the precise assignments in Figure 1 can seem rather arbitrary. Worse, even as it says too much, the paraphrase also fails to say enough. As Donald Davidson (1984, 263) puts it, "when we try to say what a metaphor 'means', we soon realize there is no end to what we want to mention." A symptom of these twin inadequacies is that so many metaphorical paraphrases "live on intimate terms with qualifiers like 'roughly' and 'and so on'" (Hills 1997, 121). But if metaphorical meaning is inherently indeterminate and inexhaustible, then it would seem that no definite, finite paraphrase, however sophisticated and carefully worked, could in principle do justice to the original.

The dissatisfaction with paraphrase articulated here is deep and often genuinely warranted.⁴

However, to the extent that the argument succeeds in undermining the paraphraseability of metaphor, it also thereby undermines the paraphraseability of many literal utterances. Suppose you asked me to restate just what I meant in saying any of the following:

(3) Jane is a real woman now.

(4) He's a politician's politician.

(5) He is an honorable and upright member of the petit bourgeoisie.

(6) She thinks she's hit the big time, living the glam LA lifestyle, but she's just another aspiring waitress.

Much as with (1) or (2), I would offer several candidate constituent clauses to my claim, highlighting some and downplaying others. I would also want to insist, though, as with (1) and especially (2), that any such restatement left out part of what I was originally after, while also forcing me to make more explicit and definite commitments than I originally had in mind. My original utterance nicely reflected my thought's own roughness and openness; the paraphrase does not.

Much ordinary talk — let alone literary writing — is loose and/or evocative in just this way, *despite* being literal. I believe that this is because much of our talk, literal and metaphorical, relies upon characterizations, whether pre-fabricated as with stereotypes, or made ready-to-order as in a narrative. Characterizations are almost always merely rough and intuitive, and go well beyond the content that's conventionally encoded in our concepts and words.⁵ Paraphrases of such utterances will often be merely partial and approximate, and will impose a determinacy on the speaker's meaning that it did not originally possess. The difficulty here, however, is a general one with paraphrasing utterances whose intended contents are not fully determinate; it is not a difficulty that applies to metaphor in particular. We don't usually conclude that literal utterances like (3) through (6) lack propositional content, or that their content is importantly different in kind from that expressed by utterances which can be paraphrased more easily.

⁴ Though not always: the content of a speaker's metaphorical speech act may be quite determinate, especially (but not exclusively) in ordinary conversational contexts. For instance, Anne Bezuidenhout (2001, 157) imagines someone being described as "a bulldozer" in a conversation about who should become department chair, where this clearly means that this person is forceful, impervious to distraction, and unlikely to be halted by obstacles.

⁵ Advocates of 'prototype' theories of concepts (e.g. Rosch 1999) are likely to reject this claim — although, once it becomes clear just how rich characterizations are, they may prefer to distinguish prototypes from what I call characterizations. In any case, I think concepts are not prototypes, for familiar reasons (cf. Fodor and Lepore 1996): for instance, prototypicality is a matter of degree, while membership in a natural kind is not.

Thus, the argument can't be used to establish the special, irreducible ineffability of metaphorical meaning.

The final argument against paraphrase implicit in the quote from Black is that the paraphrase fails to give the same "insight" as the original. The metaphor doesn't merely express a propositional thought; it reveals an overall perspective — a "filter" (Black 1962) or "frame" (Moran 1989) for organizing and coloring our thoughts, both about the subject under discussion and about the world at large. This is what we 'get' when a metaphor makes us "see things in a new light" or under a new "aspect" (Davidson 1984); and it is what people are typically after when they appeal to metaphor's fecundity and mystery.

I think that metaphors have this 'aspectual' effect because they employ one characterization to structure another. In speaking metaphorically, I believe, the speaker intends for her hearer to make his characterization of the subject under discussion as structurally isomorphic as possible, given limitations of conversational relevance, to the governing characterization indicated by the metaphorical term. In conversations where the speaker intends to make a determinate point, this merely requires identifying a few features in the subject characterization which can be matched to prominent features in the governing characterization. For richer, more 'deeply meant' metaphors, however, the speaker wants his hearer to take the project of applying the governing characterization more seriously.⁶ If Romeo had directed his soliloquy at an audience besides himself, for instance, then his utterance would have invited his hearers to use their characterization of the sun to structure their overall characterization of Juliet, and indeed their characterizations of those around her as well: of Rosaline and the other ladies of Verona, and of Romeo himself and his love for Juliet.

As essential as 'aspectual thought' is to metaphor's workings and to its overall effect, however, I do not think it belongs in the paraphrase. As I said in §1, a paraphrase should capture the content of the speaker's intended speech act. But a 'perspective' or an overall way of structuring one's thinking does not itself fix any conditions of satisfaction, and so it cannot be the content of (e.g.) a claim.

⁶ In such cases, the speaker's communicative intentions are quite open-ended, and may even include the intention to let the hearer conspire in determining the metaphor's content.

‘Perspectives’ are indeed cognitive, in the sense of being *tools* for thinking, but they are not themselves thoughts. Thus, while Black is correct both that a literal paraphrase lacks the same insight as the original, and that this is a significant cognitive loss, this does not impugn the paraphrase’s own adequacy.

Finally, notice that even this phenomenon, of inducing an openended ‘aspect’ or ‘perspective’, is not distinctive of metaphor. The discovery that Donald Rumsfeld was a wrestler, for instance, gave me a whole new schema for understanding him. Similarly, a guiding principle or proverb, like ‘People are basically selfish’, or ‘What would Jesus do?’, can structure our thinking about a wide range of information and situations. Thus, even if someone were to insist that aspectual effects should be included in a paraphrase, this still would not show that metaphor does something that literal speech cannot.

4. Demonstratives and Beyond

Rather than focusing on metaphor’s complexity, I think a more promising avenue for establishing the claim that metaphors pose a distinctive challenge to paraphrase lies in the *precision* of at least some metaphors. The fact that any given language contains only a finite number of fixed linguistic expressions constrains the range of properties its speakers can talk about directly. Demonstratives extend these linguistic resources significantly, by enabling speakers to exploit the world itself in order to construct novel expressions. Metaphors can function communicatively much like demonstratives in this respect. For example, characterizing the drunk at the bar as a “wheezing bagpipe” allows me to capture the particular tone of his voice: loud, braying, continuous, nasal. These latter adjectives provide you with a general schema for imagining the relevant sound, but the metaphor is considerably more vivid and precise, because it exploits your specific, experiential knowledge of the sound that bagpipes make.

In this case, the speaker could employ a demonstrative in place of the metaphor, by mimicking the drunk’s voice and then ostending that sound. However, it is not always possible to replace a metaphor with a demonstrative in this way. We can only construct a demonstrative if we can ostend a sample of the relevant property; but this isn’t always possible. When it’s not possible, the speaker can still presumably form an appropriate demonstrative in her own mind, because she has an experiential memory of the

relevant property. But unless her hearer also happens to be previously acquainted with that property, she will have no hope of getting him to grasp the verbalized demonstrative's content. His limited range of experience precludes him from any cognitive access to the property the speaker is thinking of.

Metaphors, however, don't require ostension, and so they can allow the speaker to communicate her intended content even when ostension is ruled out. The metaphor accomplishes this by setting up an implicit analogy between two object-property pairs, where the hearer presumably has had experience with both the object and the property in one pair but only with the object of the second. This is how a metaphor like (7) works, I think:

(7) When he finally walked out the door, I was left standing on the top of an icy mountain crag, with nothing around me but thin cold air, bare white cliffs, and a blindingly clear blue sky.

Here, the speaker is claiming to have experienced a specific property, one for which the language has no existing expression, and one which the hearer has not (let us suppose) experienced himself. Such cases are not, I think, all that unusual. In particular, in addition to emotions, we also often describe sensations in the metaphorical terms of distinct sense modalities — as a glance at wine appreciators' magazines demonstrates.⁷ In general, there is no clear algorithm for solving the analogical equations set up by such metaphors; they require a range of imaginative skills, both cognitive and empathetic. However, we typically are able to identify the speaker's intended solution to such equations. And when we do, the metaphor can increase our body of knowledge, by giving us a new concept for an unfamiliar property.

Even in this sort of case, it's important to see, the speaker herself is still not forced to *speak* metaphorically. She could communicate her intended content metalinguistically, as in

(7₁) I felt the way one might metaphorically describe as standing on top of an icy mountain crag...

Slightly more helpfully, she could make the metaphor's implicit analogy explicit, as in

(7₂) I felt an emotion which is like the physical feeling of standing on an icy mountain crag....

⁷ Another example: Björk apparently asked her choir to make their singing "like marzipan" (Alex Ross, "Björk's Saga," *The New Yorker*, 8/23/04), to which they responded, "Oh, you mean *dolcissimo*." But it seems plausible she had a more specific request in mind: that they make a specific sound analogous to the specific taste of marzipan, which combines sweetness, acrid bitterness, and unctuousness in a particular way.

These paraphrases are technically literal, and so do eliminate the metaphor *per se*. But I don't think they count as adequate paraphrases, because they still rely at least implicitly on the original metaphor, and so they fail to provide explicit formulations of the speaker's meaning. This is obvious, I take it, for (7₁); (7₂) calls for a bit more consideration. If 'like' expresses a relation which holds just in case there is *some* similarity between the compared objects (cf. Goodman 1972), then (7₂) is vacuously true, and fails to capture the speaker's intended content, which in this case is quite specific and determinate. On the other hand, if 'like' expresses a substantive relation which holds just in case a particular, contextually salient similarity holds between the two objects (cf. Tversky 1977), then (7₂) implicitly builds those similarities into its content. It may then succeed in capturing the speaker's intended content,⁸ but it arguably also fails to be fully explicit, in much the way that 'He's ready' fails to specify its implicit argument.

We could address this latter worry by citing the relevant similarities explicitly, as in

(7₃) I felt an emotion which is like the physical feeling of standing on an icy mountain crag... in respects *i, j, k...*

But this renders another difficulty, already present with (7₂), more apparent: construed as a paraphrase, (7₃) attributes unintended content to the speaker. In uttering (7), the speaker isn't making any claims about what icy mountain crags are like or about their relation to her emotional state — she's just characterizing her emotion, using shared attitudes about icy mountain crags to do so. Her intended claim has the form: "When he left, I felt *that way*". A paraphrase like (7₃), or even (7₂), inappropriately builds facts about icy mountain crags, and the relations between the relevant emotion and the physical feeling of standing on icy mountain crags, into the content of the speaker's claim. These facts may be part of the implicit background conditions which enable the speaker to employ *this* metaphor as a vehicle for communicating *that* content, but they aren't actually part of that content itself.

It thus seems that any paraphrase of a demonstrative-like metaphor like (7) will be torn between two conflicting criteria: providing an appropriately literal and explicit specification of the speaker's

⁸ There's a further issue about whether, on a substantive view of 'like,' (7₂) would be true or false, depending on whether one takes the similarities relevant for interpreting the metaphor or its corresponding simile to be contextually salient. Many similes seem to be literally false, because they rely on similarities which are not independently salient even within their context of utterance; see Fogelin (1988).

meaning, and respecting the simple, singular content of the speaker's intended claim. It may appear that we can meet both these criteria if we prefix the specifying material with 'Dthat':⁹

(7₄) I experienced Dthat[the emotion which is like the physical feeling of standing on an icy mountain crag... in respects *i, j, k...*]

But I think this too fails. I said in §1 that an adequate paraphrase should consist in a set of sentences whose overall semantic content is the same as the content of the speaker's intended speech act. And of course, this semantic content must itself depend upon the semantic values of the sentences' constituent parts. In particular, the content of all the candidate paraphrases must depend upon which physical feeling is in fact denoted by "the physical feeling of standing on an icy mountain crag...". Pragmatic interpretation, by contrast, merely requires speakers and hearers to coordinate interpretive assumptions so as to converge upon the appropriate content; it doesn't require that those assumptions be true, or even that the speakers' and hearers' reasoning about those assumptions be correct. Metaphorical interpretation, in particular, calls for speakers and hearers to coordinate their *characterizations* of the relevant phenomena; and characterizations need not be accurate, or even believed to be accurate.¹⁰ Thus, because pragmatic interpretation in general, and characterizations in particular, needn't proceed by way of how the world really (possibly) is, while semantic content necessarily does so proceed, a paraphrase like (7₄) won't necessarily deliver the same content as the speaker's intended claim.¹¹

⁹ Presupposition is another mechanism for narrowing the content claimed without building extra content into that claim itself. However, because the presupposed material is playing an essential role in fixing the truth-conditions for the paraphrase, it should still count as part of the paraphrase in its entirety. Thus, presupposition is at best effectively equivalent to 'Dthat' in this respect, and the worry about 'Dthat' applies to presupposition as well: a paraphrase which includes presuppositions that appeal to merely pragmatic assumptions shouldn't count as adequately literal and explicit. Further, we still need to appeal to 'Dthat' in order to respect the singular nature of the speaker's claim. Much the same objection applies to an introductory imaginative exercise which attempts to induce the relevant feeling directly in the hearer (e.g., 'Imagine that you're on top of an icy mountain crag...'): the description of the imaginative exercise itself contributes to fixing the truth-conditions, and so should be included in the paraphrase. Thanks to Thony Gilles for pressing the possibility of presupposition.

¹⁰ Suppose, for instance, that the speaker and hearer's characterizations of icy mountain crags were derived entirely from Jack London novels and other adventure writing, but that the depictions of icy mountain crags in that writing was false. Suppose even that the speaker and hearer both knew those depictions to be inaccurate. This would affect neither their coordinated interpretation of the metaphor, nor the hearer's evaluation of its truth.

¹¹ What about a paraphrase like 'Dthat [the emotion the speaker is thinking of]', or 'Dthat[the emotion the speaker experienced 12 years ago]'? These do get the content right, but in a way that renders the project of paraphrasing utterances trivial in general. **In particular (cf. §5 below), such a paraphrase fails to enable a speaker to introduce a new term denoting the intended property into the common language, because such a paraphrase wouldn't enable a hearer to reapply the relevant term to new instances of the intended property.**

It might now seem that we can side-step this difficulty by building characterizations themselves into the paraphrase. But characterizations are paradigmatically pragmatic beasts, and the process of matching similar features between characterizations essentially involves leaps of logic and imagination. Because these phenomena lack the stable, broadly rule-governed behavior of semantic meaning, including them in the paraphrase would undermine the paraphrase's adequacy: an otherwise linguistically competent, rational hearer could no longer understand the paraphrasing sentence simply in virtue of his basic linguistic competence and rationality. To understand the paraphrase, he would need to construct a contextually appropriate characterization of icy mountain crags, exploit the mentioned similarities to construct a characterization of the appropriate emotion, and then imagine how that emotion would feel, all the while guided by his sense of what the speaker might be after. That is, he would need to engage in just the sort of interpretation called for by the original metaphor.

5. The Principle of Expressibility

Unless adequate, literal, explicit paraphrases can be constructed for metaphors like (7), such cases will count as counterexamples to what John Searle has called the "Principle of Expressibility." Searle's Principle serves as a nice foil for our discussion, I think, because it encapsulates the rationale that leads many philosophers to assume that all metaphors must ultimately be paraphraseable.¹² Searle's statement of the Principle is as follows:

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 (1969, 20).

It's important to be clear on the role played by the possibility operator here. Searle allows that a language may not in fact "contain the words or other devices for saying what I mean" (1969, 19); he even points to metaphor as proof of this point (1979, 114). Rather, the central claim is that when there is a semantic

¹² Davidson also seems to appeal to something like the Principle of Expressibility in arguing from the failure of paraphraseability to the conclusion that metaphors lack any cognitive content. He writes (1984, 260): "If a metaphor has a special cognitive content, why should it be so difficult or impossible to set it out?...Can't we, if we are clever enough, come as close as we please?"

inadequacy, then “I can in principle at least enrich the language by introducing new terms or other devices into it” (1969, 20). That is, there is nothing about the content itself that a speaker might have in mind which could prevent its being the semantic value of a linguistic expression.

In a case like (7), however, the speaker lacks the resources necessary to introduce a new term in the manner Searle imagines, even if she avails herself of all possible literal means. She herself could easily invent a new word, relying on demonstrative reference through memory. But she could not introduce that word into the language, because her hearer would be in no position to comprehend it — not as a result of linguistic incompetence, or irrationality, but just from a lack of worldly experience.

Searle tries to fend off the trouble caused by cases of this general sort by denying that the Principle applies to expressibility in a public language. He writes,

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. (1969, 19)

That is, Searle reads the possibility operator in his statement of the Principle as requiring only that there be *some* possible circumstance in which an expression *E would* express the speaker’s meaning — a condition that clearly is met in the case of (7). However, this reading is unacceptable, because it weakens the Principle to the point of theoretical irrelevance. Searle claims that the Principle of Expressibility “has the consequence that cases where the speaker does not say exactly what he means — the principal kinds of cases of which are nonliteralness, vagueness, ambiguity, and incompleteness — are not theoretically essential to linguistic *communication*” (1969, 21, emphasis added). But metaphor, and nonliteralness more generally, can be dismissed as theoretically inessential to communication only if alternative means of communication are in principle available. This conclusion fails to follow from the weakened Principle, which no longer requires that what is meant can be said in terms that one’s audience can even in principle understand, given cognitive resources sufficient for linguistic competence.

It is of course true that *after* the speaker has gotten her hearer to identify the relevant property by metaphorical means, she can then introduce a new term which denotes it. Linguistic change often occurs

in just this way, whether through explicit definition or through slow conventionalization. But because the metaphor here plays an essential role in defining the new term, this possibility cannot be used to show that metaphor in general is theoretically eliminable. Although each particular metaphor can eventually be eliminated, the situation exemplified by (7) can always arise anew for a different property. And so long as this is possible, there will be metaphors which cannot be paraphrased in fully literal, explicit terms.

6. Metaphors in Thought

Although Searle's Principle can't be used to establish his desired conclusion, I do think there is a natural inclination to dismiss cases like (7) on something like the grounds he cites. That is, the impossibility of literal expression here clearly depends upon merely contingent factors about communication, resulting from differences in our limited, only partially overlapping bodies of experience. It doesn't follow from ineluctable facts about content, intentionality, reference, and the like. On the one hand, a different language — French, say — may contain a fixed linguistic expression referring to a property that we can only express metaphorically in English; linguistic competence with French would then provide cognitive access to the relevant property. On the other hand, given an appropriate context of utterance, the speaker could ostend the property directly, and her hearer would be able to understand her. Given this, the barrier here to literal expression may not yet seem very deep or interesting. To allay this dismissive reaction, I want to push the argument one step further, from communication to cognition. Not everything that can be meant can necessarily be given literal expression, *even* in a private language.

Suppose, if you will, that I am investigating a relatively ill-understood realm of phenomena to which I have only indirect experiential access: for instance, sub-personal cognitive processing. As my research progresses, I begin to believe that there's a certain causally efficacious property which I want to focus on. How can I identify this property? I don't yet have a clear fix on just what its causal powers are, and so I can't define it as 'the property that plays causal role *R*'. I also can't identify it or its effects ostensively — as '*this* property' — because I only ever encounter it embedded within a complex field of interacting processes. Not only can't I observe it itself directly and in isolation; I also lack any reliable

way to distinguish its role in producing a given effect from the roles simultaneously played by a host of other properties. In such a circumstance, though, a metaphor may enable me to ‘lock on’ to the particular property I’m interested in, much as (7) enabled its hearer to grasp the property its speaker had in mind.

Richard Boyd (1979) argued 25 years ago that this was precisely the situation faced by researchers in cognitive science; his claim still seems true today. If one is any sort of functionalist about cognition, then one believes that some properties of cognition are essentially individuated by their causal relations. We have a rough idea about what some of these properties’ causal relations must be, and we know considerably more about their effects, but we can’t yet define the properties in fully literal terms. We can still make theoretical and experimental progress, though, by thinking metaphorically — for example, by exploiting the metaphor of memory storage and retrieval as the opening of a computer file (where this is itself, of course, a metaphor drawn from physical data storage). Research progresses, in part, by investigating specific candidate similarities that might underwrite the analogical equations that are implicit in such metaphors. As we establish some similarities and rule out others, our cognitive access to the properties under investigation becomes more fully and literally conceptualized.¹³ At some point, if investigation progresses well, we may well be able to dispense with the metaphor in favor of a new, literally applicable concept. But at this early stage of our inquiry, the metaphor plays an essential role in fixing what we are thinking about.

I can see two main ways to resist the claim that metaphor is essential here. First, just as with (7), one might insist that the metaphor *per se* can be eliminated. I said that the metaphor fixes the property *x* which I want to investigate *via* an implicit analogical equation. But we can make this equation explicit in a literal definition, as for example in:

¹³ It’s important to remember that investigation involves both the ruling *in* and the ruling *out* of relevant similarities, and hence of possible denotations. Progress often requires precisely delimiting the scope of the metaphor. Thus, for instance, as biologists’ understanding of genetic replication has become more nuanced, it has become clear that only some of the applications which have been suggested for the metaphor of genetic ‘coding’ are substantive. Peter Godfrey-Smith (2002) surveys the history of this metaphor’s application; he argues that because its current use is inappropriately broad, biologists fail both to grasp the robustness of its genuinely substantive application and to consider alternative models for phenomena for which the metaphor is not useful.

- (8) The property of cognition that causes memory retrieval in a manner that is analogous in some theoretically relevant respect to opening a folder in a computer program.

While the metaphor is considerably more compact and convenient than (8), the thought itself does not require that form of expression. However, as I argued in discussing (7), identifying the denotation of a literal description like (8) requires the same cognitive capacity as the original metaphor does. We still need to identify which particular similarities are relevant, and then construct a positive concept of the appropriate property on that basis. And here, unlike with (7₃), say, no one is currently in a position, even in principle, to identify just what the relevant similarities are. Further, in solving the equation, we don't simply focus on the source feature in isolation: say, on which actions are involved in opening a computer program's file. We also exploit our knowledge — and specifically, our know-how — of related actions, objects, and properties: of the uses to which we put files, of the relations between files within the same folder, of how we decide what belongs in the same file and what deserves its own file. As Nelson Goodman emphasized, in metaphor “a whole apparatus of organization takes over new territory” (1968, 73), and this organizational apparatus is still operative in making sense of a literal description like (8). Thus, just as with communication, we can eliminate the metaphor *per se*. But the general analogical patterns of thought and the specific surrounding background knowledge on which metaphorical comprehension relies still remain essential to the comprehension of its literal substitute.

This emphasis on how much work remains even after formulating the appropriate analogical equation leads to the second objection, though: why think that an unsolved equation of this sort can give us cognitive access to a specific feature in the world? Just as a broad, sweeping gesture toward a crowd fails to isolate a particular individual as the referent of ‘that guy’, so, it seems, both the metaphor and its literal restatement fail to isolate any particular property as their referent. Worse, **unlike in the communicative case, here there are no specific speaker's intentions which could supply the needed determinacy**. Our task here isn't just to recover the speaker's intended referent, where *she* knows perfectly well what she's talking about, but rather to reach out to a dimly glimpsed feature in the world.

I think this objection seems damning only so long as we remain at a quite abstract level. In the context of actual investigation, our general theoretical assumptions and interests do impose substantive constraints on the metaphor's interpretation. On the one hand, not just any similarity between elements in the two domains is an acceptable candidate for solving the equation. In this case, the background assumption that minds are fundamentally concerned with information processing significantly constrains the respects in which we are willing to count memory retrieval as like a software program's opening a file. On the other hand, the domain of properties which could possibly count as solutions to the analogical equation, and so as referents of the metaphor, is heavily restricted by independent assumptions about the general field under investigation: here, by assumptions about the mechanics of neural processing. So, although there must be some significant indeterminacy — otherwise, we could dispense with the metaphor after all — I think these constraints on the range of possibly relevant similarities and on the domain of possible referents are sufficiently substantive to fix genuine conditions of satisfaction for the metaphor, despite the fact that we ourselves don't fully understand what those conditions are.

A more plausible version of the objection is that the metaphor's interpretive indeterminacy threatens to introduce referential vagueness as well. Perhaps several distinct cognitive mechanisms are each like opening a file in a different theoretically relevant respect, so that they all seem like equally plausible referents of the metaphor. However, while this is indeed a live possibility, it potentially plagues any attempt to talk about something with which we are not fully and intimately acquainted, even when our modes of reference are entirely literal. Scientific progress often consists in discovering that a given term is vague, or has multiple partial denotations, and then altering its use on that basis; consider here the history of investigations into heat, or mass, or Vulcan.

It may be useful in this context to invoke Descartes on our acquaintance with the nature of God:

It is possible to know that God is infinite and all powerful although our soul, being finite, cannot grasp or conceive him. In the same way we can touch a mountain with our hands but we cannot put our arms around it as we could around a tree or something else not too large for them. To grasp something is to embrace it in one's thought; to know something, it is sufficient to touch it with one's thought. (Letter to Mersenne, 27 May 1630 (AT 152)).

I suggest that more of our cognitive engagement with things is dim and groping, at least in our initial dealings with them, than philosophers usually like to admit. Unlike the experiential properties picked out by a metaphor like (7), we will never be able to fully “embrace”, or be acquainted with, the sorts of features I’ve been talking about here: causal properties embedded in a complex field of interactions, with which we cannot in principle interact directly. But we can still “touch” such features: we can feel the effects of, think about, investigate, and even have knowledge about them. Metaphor provides one important initial means of such cognitive access, even if the aim of scientific investigation is eventually to render the metaphor itself disposable.

7. Conclusion

As a defense of metaphor, this may all seem rather disappointing. I haven’t identified a general, distinctive, irreducible and ineffable feature of metaphor, which is what partisans of metaphor have typically been seeking. But both metaphorical and literal talk are too various for us to expect such a feature. Philosophers have repeatedly insisted, both in defense and in dismissal of metaphor, that it is deeply different from literal communication because it is indeterminate, complex, rich, evocative, and openended. They have failed to notice that not all metaphors fit this model, and that much literal speech does. Thinking clearly about what metaphor does, and how it does it, forces us to acknowledge just how much of our communicative and cognitive lives exploit the same patterns and processes of thought that we find in metaphor. This is not to say that these cognitive modes and abilities are essential to thought or to talk *per se*. I have offered no reason to think that beings who did not employ characterizations, aspects, and analogies would be incapable of representing the world to themselves and to each other. Rather, my claim has been, first, that these patterns of thought play an important role in our distinctively human form of engagement with the world; and second, that in certain circumstances exploiting them can genuinely extend our basic cognitive and communicative resources.

Appendix: Romeo's speech on Juliet, II.ii.1ff.

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.

References

- Bezuidenhout, Anne (2001): "Metaphor and What Is Said: A Defense of a Direct Expression View of Metaphor," *Midwest Studies in Philosophy* 25, 156-186.
- Black, Max (1962): "Metaphor," in *Models and Metaphors: Studies in the Philosophy of Language* (Ithaca: Cornell University Press).
- Boyd, Richard (1979): "Metaphor and Theory Change: What is 'Metaphor' a Metaphor For?," in A. Ortony (ed.), *Metaphor and Thought* (Cambridge: Cambridge University Press).
- Brooks, Cleanth (1947): "The Heresy of Paraphrase," in *The Well Wrought Urn* (New York: Harcourt Brace).
- Davidson, Donald (1984): "What Metaphors Mean," in *Inquiries into Truth and Interpretation* (Oxford: Clarendon Press).
- Fodor, Jerry and Ernie Lepore (1996): "The Red Herring and the Pet Fish: Why Concepts Still Can't be Prototypes," *Cognition* 58:2, 253-270.
- Fogelin, Robert (1988): *Figuratively Speaking* (New Haven: Yale University Press).
- Glucksberg, Samuel, and Keysar, Boaz (1993): "How metaphors work," in A. Ortony (ed.), *Metaphor and Thought*, 2nd ed. (Cambridge: Cambridge University Press).
- Godfrey-Smith, Peter (2002): "On Genetic Information and Genetic Coding," in P. Gardenfors, J. Wolenski and K. Kajania-Placek (eds.), *In the Scope of Logic, Methodology, and the Philosophy of Science*, vol. II (Dordrecht: Kluwer).
- Goodman, Nelson (1968): *Languages of Art: An Approach to a Theory of Symbols* (Indianapolis: Bobbs-Merrill Company).
- (1972): "Seven Strictures on Similarity," in *Problems and Projects* (Indianapolis: Bobbs-Merrill).
- Grice, Paul (1989): "Logic and Conversation," in *Studies in the Way of Words* (Cambridge, MA: Harvard University Press).
- Hills, David (1997): "Aptness and Truth in Verbal Metaphor," *Philosophical Topics* 25:1, 117-153.
- Levinson, Stephen (2000): *Presumptive Meanings: The Theory of Generalized Conversational Implicature* (Cambridge MA: MIT Press).
- Moran, Richard (1989): "Seeing and Believing: Metaphor, Image, and Force," *Critical Inquiry* 16, pp. 87-112.
- Nogales, Patti (1999): *Metaphorically Speaking* (Stanford, CA: CSLI Publications [Center for the Study of Language and Information]).
- Rosch, Eleanor (1999): "Principles of Categorization," in E. Margolis and S. Laurence (eds.), *Concepts: Core Readings*, (Cambridge MA: MIT Press).
- Searle, John (1979): "Metaphor," in *Expression and Meaning: Studies in the Theory of Speech Acts* (Cambridge: Cambridge University Press).
- (1969): *Speech Acts: An Essay in the Philosophy of Language* (Cambridge: Cambridge University Press).
- Tversky, Amos (1977): "Features of Similarity," *Psychological Review* 84:4, 327-352.