

Against Metaphorical Meaning

Ernest Lepore and Matthew Stone¹

Speakers can use metaphors to make their views plain. Matt Groening, with (1), presents an unmistakable characterization of the feelings of love.

1. Love is a snowmobile racing across the tundra and then suddenly it flips over, pinning you underneath. At night, the ice weasels come.

Groening's imagery has an absurd particularity, of course, yet many readers still resonate to his portrayal of the exhilaration of a new relationship, and in particular, to the gnawing, trapped, painful feelings that too often appear once the exhilaration subsides. Suppose we grant that love really does follow a trajectory from excitement, to misadventure, to torment. Is that what Groening *means* by (1)? Is he, thereby, using (1) to *communicate something true*?

It is tempting to answer these questions in the affirmative. Once we do, we can then describe metaphorical interpretation in the traditional semantic and pragmatic terms, thus, assimilating metaphor to other familiar uses of language. Moreover, there are a range of familiar theoretical tools to allow us to implement this view. Perhaps, metaphors supply new, context-dependent meanings to words and phrases (Stern, 2000); or enrich and adapt our ordinary meanings to more relevant alternatives (Carston, 1991; Bezuidenhout, 2001; Recanati, 2001). Maybe their uses carry implicatures beyond literal meaning (Grice, 1989; Searle, 1979); or let us construct propositions by distinctive metaphorical thinking, substituting for the ordinary meaning in further steps of semantic and pragmatic interpretation (Black, 1955; Hills, 1997; Camp, 2003). In each of these cases, metaphorical usage carries a meaning its speaker is trying to communicate.

In this paper, we will reject this commonplace view, and propose in its place the view that, though metaphors can issue in distinctive cognitive and discourse effects, they do so *without* issuing in metaphorical meaning and truth, and so, *without* metaphorical communication. Our inspiration derives from Davidson's critical arguments against metaphorical meaning (1979) and Rorty's exploration of the diverse uses of language (1987). But unlike these authors we ground our discussion squarely in distinctions about causal mechanisms in cooperative activity developed by Grice (1957, 1969) and others.

Language figures in many cooperative activities where speakers are understood only when their intentions are recognized, but not all such exchanges constitute communication, at least not as conceived by Grice, and so, not all of them involve speaker meaning: for example, our practices of hinting, joking, trash-talking, flirting, and flattering, among others, involve crucially linguistic behavior; each normally includes

¹ We are grateful to Sam Cumming, Alex Lascarides and Dan Sperber for feedback on this work. This work was supported in part by NSF grants CCF-0541185 and HSD-0624191.

linguistic communication as traditionally understood, but essential aspects of these practices cannot be reduced to, or characterized in terms of, the communicative enterprise. In what follows, we will elaborate on these vague claims as well as argue that metaphor is best catalogued along with these practices and *not* with communicated meaning. This position on metaphor is compatible with a wide variety of positive proposals about the processes of metaphorical understanding and their cognitive effects, any one of which affords appealing responses to commonly-voiced criticisms of the Davidsonian idea that metaphors lack meaning.

1 Metaphor and Communication

We are unaware of any argument for or against metaphorical meaning and truth that includes a discussion of speaker meaning. This is surprising, since anyone who claims metaphorical usage carries meaning, in one manner or another, say, to underwrite assertion, truth, disagreement, or interpretation, must be assuming speakers who employ metaphors mean something by them, something that exceeds their literal meaning. Yet the standard progression from Aristotle forward has been to explain how to assimilate metaphorical meaning to standard semantic and/or pragmatic accounts and mechanisms. All such efforts are premature until it has been established that whatever point a metaphorist intends to make with his usage rises to the level of speaker meaning. To discern whether it does, we will adopt the familiar perspective that speaker meaning *is* a species of *intention in action*.

The connection between speaker meaning and intention traces back to the landmark discussion of Grice (1957). What distinguishes speaker meaning from other intentions an agent might have is the kinds of effects his action aims for, and how these effects are to be achieved. Grice describes the distinctive effects of meaning in terms of eliciting a decision in audience members that is under their control. As canonical examples, Grice cites inducing an audience to adopt a belief, in the case of an assertion, and inducing an audience to commit to an action, in the case of an imperative. In these cases, if *p* is the content of the belief that the audience adopts, or a description of the action that the audience commits to, then we can say *p* gives the intended effect the speaker is aiming for.

Grice captures the distinctive process of communication by requiring that any intended effects be elicited in part by an audience's recognition of the speaker's intentions. In other words, in order for a speaker's intention to *count* as what his utterance means, this intention must envisage that its recognition will provide an audience with the decisive reason to commit to the very effect the speaker is aiming for. To take a concrete illustration, in order for a speaker to mean that it is raining with an utterance of "It is raining," he must intend that his audience believe that it is raining; and further that they recognize that the speaker has this intention that they should believe that it is raining; and finally, that it be by means of this very recognition that they actually come to believe that it is raining. Generalizing, then, Grice's explication of speaker meaning comes to (2):

2. In order for a speaker *S* to (speaker) mean *p* with an utterance *u* is for *S* to produce *u* with the intention to elicit the effect given by *p* in an audience by means of their recognizing that this is his intention (cf. Grice, 1989, p. 220).

(2) locates speaker meaning squarely in how an agent's actions impact on others. When an audience considers meaning in utterances, they are, according to Grice, asking what was the speaker aiming for and how did he envision it would transpire.

Grice's aim, in part, is to distinguish *meaningful signaling* from other ways in which we share information with one another. On the one hand, Grice is concerned with explaining (away) cases where actions clearly provide us with information but where we fail to locate any speaker meaning in them. To take Grice's vivid example, Herod's action of presenting Salome with the head of John the Baptist clearly informs her that John is dead; but it would be a mistake to conclude Herod's action means that John is dead; his *act* of showing her the head doesn't signal this information; it's rather seeing John's severed head that leads her to believe he is dead. It is true Herod intends Salome to learn that John was dead, and she knows this, but, again, she learns he is dead from seeing his severed head, *not* from recognizing Herod's intention.

This doesn't mean language is essential to meaning. Grice sees meaning in language as part of a larger phenomenon of communication, and not just something that depends distinctively on the use of linguistic forms or on the existence of linguistic conventions or grammar. (Grice would actually prefer to reduce our knowledge of meaning in language to our knowledge of psychology: what a sentence means, in his view, is just what speakers generally mean when they use an utterance of the sentence.) This is why his analysis is formulated in terms of fundamental notions from collaboration and *not* from linguistic theory. For whenever we participate in joint activities, we intend our contributions to be recognized; often we go further and rely on specific processes and mechanisms in doing our part.

Information can be contributed in ways besides through language, and many of these ways depend on just the sort of coordination Grice associates with intention recognition. Grice cites as an example the information retrieved from a bell on a tram. When the conductor rings it three times, this means the tram is full. Lewis (1969) considers the signals from the ride of Paul Revere. When the sexton of the Old North Church hangs one lantern in the tower, this means the Redcoats are coming by land. In both cases, an agent intends to convey information to an audience; and in both, the audience must recognize the relevant intention in order to tap into the information the agent is contributing.

In this paper, we build on the distinctions first targeted by Grice, and evaluate whether metaphor is best classified in this light as generating speaker meaning. The details as Grice lays them out in (2) are tendentious; the literature offers a number of difficult cases. This requires us to clarify the status of these alleged counter-examples before we can settle questions about metaphorical meaning sharply. Broadly speaking, three responses are possible. Grice might be right: properly understood, the notions of deciding on an effect and acting on the basis of a recognized intention dissolve these alleged counter-

examples. If this is your view, you can push our arguments through with Grice's original characterization (2). Alternatively, you might think Grice's view needs tweaking. Grice himself issued a range of emendations and qualifications up to the William James Lectures, published as (1989, p.114), giving secondary conditions on the kinds of intentions that count as issuing in speaker meaning and the way in which those intentions must be recognized and used. His final account involves some six or seven bullet items, and takes almost a page to state; it is not particularly natural. Nevertheless, if you endorse it, you will be able to push our arguments through. And, lastly, you might want to precisify (2) invoking specialized constructs from pragmatics; in particular, following Lewis (1969, 1979) and Thomason (1990), by giving explicit and direct theoretical characterizations of the notions of coordination and of contributing information to an interaction. These constructs don't necessarily make for a simple or elegant philosophical analysis, and they undercut some of the metaphysical aims under which Grice proposed (2). For that reason, they may prove more congenial to linguists or psychologists interested in the status of metaphor in the science of communication. But if you agree with them, you will enable you to push our position on metaphor through. To see what we have in mind, consider a counter-example to (2) that concerns the aims of meaning.

When a student answers a question on an oral exam, she surely means what she says. But how can we square this intuition with Grice's analysis? If her answer is correct, mustn't the examiner already have this information? But if her aim is for the examiner to credit her with a correct answer, isn't his response a function of the standards of the exam, *not* of his recognition of her intention to be judged correct? How does her answer contribute this information to the interaction? Such cases highlight a fundamental difficulty in characterizing the effects of meaning concretely in terms of decision or belief, or even, following Grice (1969), generally, in terms of the production of a suitable response by the audience. In this regard Lewis's (1979) account fares better. He recommends that utterances be characterized as providing *updates to the conversational record*, an abstract 'scoreboard' that tracks interlocutors' contributions to their interaction.

Speaker meaning in this framework still involves an intention, one to update the conversational record. This record describes actions taken overtly by the interlocutors as contributions to the interaction as well as the evident, public effects those actions engender. This permits us to think precisely about what goes on the record, and to connect the conversational record to the mental states of interlocutors. Normally, when interlocutors are well-informed and attentive, they track the record, and so, come to mutually suppose, for the purposes of the conversation, that things are as recorded (Stalnaker, 1978; Thomason, 1990).²

² Stalnaker and Thomason in fact propose to reduce the conversational record to the mutual attitudes of interlocutors. However, one can also locate the conversational record by reference to the norms which interlocutors subscribe to in participating in conversation; this lets us narrate public actions as taking the same effects on the scoreboard even in defective cases where interlocutors may fail to understand one another, because they are ignorant about the facts of meaning or must work through temporary misunderstanding or ambiguity (Steedman and M. Stone, 2006; Lepore and M. Stone 2007; DeVault 2008).

Since invoking the conversational record gives an independent characterization of the effects speaker meaning can aim for, we can solve the puzzles about the presence and limits of meaning in cases such as examinations. In that case, the conversational record for the examiner and his student simply tracks questions asked and responses given; in this framework, the interlocutors' activity begins with a blank slate, not with shared background knowledge and present beliefs. The student's utterances are intended to register specific content on the record; each utterance seeks publicly to establish a specific proposition as the official answer to a stated question. Establishing an answer in this sense arises through the pure coordination of student and examiner, through the kind of causality Grice's definition aims to pinpoint. Putting the answer on the record is separate from any auxiliary responses the examiner must make (and typically, must register publicly in turn, through a suitable overt acknowledgment) to judge the answer and evaluate the student.

A complementary refinement involves Grice's account of the processes involved in meaning making. Difficult cases are not distinguished sharply by the kinds of information or reasoning that interlocutors use to understand one another. They are more clearly distinguished by the circumstances that prevail in the world, and set the stage for interlocutors' joint activity. Accordingly, we may choose to be more precise about the processes involved in speaker meaning not just in terms of intention recognition, or the interrelated attitudes of interlocutors, as in (2) and preserved through Grice (1969), but directly in terms of the kinds of situations where agents face signaling problems, following Lewis (1969). The key notion is *coordination*. Coordination problems are situations where multiple agents must match their behavior to one another, and multiple ways of doing this are possible. To act successfully, agents must use their expectations about one another to select one solution to the exclusion of others.

Updating the conversational record, so as to register a specific proposition as part of the presumptions of the interaction, is a coordination problem in Lewis's sense. There is no inherent connection between utterances and contents, so a speaker could use any utterance to register any proposition, so long as her audience used a corresponding strategy to track the propositions so contributed. Interlocutors solve the coordination problem by adhering to what Lewis calls a *convention*, that is, a particular scheme, implicit in interlocutors' mutual expectations about how they will speak and how they will understand, for aligning their efforts to track the information each contributes.

More precisely, with acting in a coordination problem, each agent adheres to a *strategy* that maps situations to chosen actions and desired outcomes. This strategy is, in effect, an intention. In fact, it is one possible intention among many. Other agents respond according to their expectations about the strategy their partners adhere to. That is, they recognize their partners' intentions and choose a response that fits what their partner has manifested – in a situation where multiple alternatives are possible. Equilibrium demands that neither agent would deviate from his decision making given more information about his partner's strategy. Focusing on coordination problems in Lewis's sense thus gives a rationale for Grice's analysis of meaning in terms of intention. Nevertheless, characterizing coordination problems directly allows us to see meaning as arising fundamentally from the *general kind* of decision making interlocutors are doing, and

general properties of circumstances in which interlocutors act, not – as in Grice’s definition – from the specific psychological processes that lead interlocutors to particular decisions in particular cases. This gives much sharper grounds to delimit coordination, and, in turn, meaning.

To concretize these various ideas, consider this case from Lewis. It is a convention that we drive on the right hand side of the road, or indeed, that we walk on the right hand side in a narrow, busy corridor. In choosing to go right, we are acting on our expectations of how others will go. We think they will also go right. If we thought they would go left, we would go left too; we wish to avoid congestion or collision. The coordinated expectations that send us all to the right – as we do in the U.S. – do not refer to anything in the world that settles our preferences or strategies for us. There is, in fact, another equilibrium. We could use the left – as in Britain or Japan. Either strategy serves equally well, so long as all agree.

Institutions and circumstances can break the symmetry which constitutes coordination. This changes the status of our decisions. In the U.S., there are sanctions for driving on the wrong side of the road. Lewis judges – and we agree – that we drive how we do because of our expectations about what others will do, given that we wish to avoid accidents. We do not drive on the right to avoid tickets. But if, in fact, we acted based on a pure preference to drive on the right, irrespective of what others did, because of the sanctions in force, then even if that strategy happened to result in an efficient solution to the problems of joint action, we would not have a problem of coordination, or a convention, in Lewis’s sense. We would act efficiently, but independently.

To take a less extreme example, many long, narrow busy corridors are fitted with two moving sidewalks, one headed toward each end of the corridor. We have a preference to take the sidewalk that goes in our direction of travel. Those with the opposite direction have the opposite preference. We all avoid collision and congestion by riding our respective sidewalks, but we do not coordinate. The world settles our preferences and strategies for us. In anticipating others’ strategies we need only assume that they act in their interests, irrespective of any consideration they give to their interactions with us, and their expectations about us. They can reason likewise. In thinking about such cases, it may be easier to apply Lewis’s notion of coordination to the circumstances that prevail in the world, rather than to work with Grice’s definition (2) and try to judge directly whether agents act on the basis of recognizing one another’s intentions.

Applied to language, paralleling the example above in discussing Grice, in order that a speaker mean *that it is raining* with an utterance of "It is raining," roughly, she must intend to contribute *that it is raining* by *coordination*. This requires her audience and she to add *that it is raining* to the conversational record in order to satisfy their joint interests, including their interest in agreeing on the record; and that her audience will do so in virtue of their mutual expectation that if she says "It is raining," then both parties agree on the record only if they all add the information *that it is raining* to it; and lastly, had different expectations prevailed and she had followed a different strategy, she could have instead done something else, thereby leading her and her audience to add *that it is raining*

to the record, and that would thereby have served their joint interests just as well. Generalizing, we can summarize with (3):

3. For a speaker *S* to mean *p* with an utterance *u* is for *S* to utter *u* with the intention of coordinating with an audience to add *p* to their conversational record.

The key difference between (2) and (3) is that instead of looking to analyze when and how recognition forms the basis for a belief, this account looks to what it means to add something to the conversational record and what kinds of circumstances and outcomes describe interlocutors' activities.

We now have a sense of what's involved in applying Grice's analysis of speaker meaning as intention more precisely and systematically. With this in place, we can revisit the key virtue of definitions like (2) and (3) for the present argument. These definitions of speaker meaning enable us to see why we do *not* take everything we can infer from someone's utterance as speaker meant.³ Two specific kinds of cases are worth highlighting: first, cases where interlocutors contribute information (by establishing a suitable belief or updating the record), but do not coordinate; and second, cases where interlocutors do coordinate (by recognizing one another's intentions or drawing appropriately on mutual expectations), but their coordination does not aim to convey information.

First, we must distinguish what Stalnaker (1978) calls "secondary effects" of utterances from speaker meaning. We can learn much from utterances beyond what they mean. Here is a simple case: when a speaker utters *u*, we learn that *she* is the speaker of *u*. But she does *not* speaker mean that she is the speaker of *u*. The information clearly becomes public on the basis of the utterance itself, and so the audience comes to believe it; it goes on the record. Interlocutors must track this information; interlocutors need it to interpret, for example, tokens of the word *I*, which refer to the speaker of the utterance. But interlocutors do not *coordinate* on who the speaker of *u* is. When *u* occurs, facts in the world, not intention recognition or interlocutors' expectations about one another, determine who said *u*.

Similar considerations establish that when a speaker utters *u*, we learn that *u* was made, but the speaker does not speaker mean that *u* was made. Interlocutors register what sounds have been produced, and how those sounds are organized into words and phrases. This is clearly public information interlocutors must track; they need it to interpret clarifications that refer to what has been said, to interpret locutions like *the former* and *the latter* that evoke expressions used.⁴ But interlocutors do not coordinate the exchange of this information. Interlocutors *may* coordinate on the performances they use to realize specific words: as with *tomato* and *tomahto*, different speakers, by convention, say the same word in different ways. But, given that the word has been used, there is no *further* problem of coordination about how interlocutors will take stock of the fact that the word

³ This was, of course, a large part of Grice's motivation in introducing (2).

⁴ And even, if Stalnaker (1998) is right, to follow the referential mechanisms that underpin the coherence of discourse generally.

has been used. The speaker has said something specific, and the fact of this utterance is manifest, so the audience comes to believe that the utterance has been made; it goes on the record. But no intention recognition or mutual expectations are involved. It is not a convention that you register *this utterance*, instead of another. Inevitably, a record of what has happened can include only what has happened.

This forces our attention to a second kind of case – where interlocutors coordinate on actions but not on information. Consider Alice and Bob, who are shopping for a dinner that they will cook together. They agreed that Alice will buy food and Bob wine. They have a coordination problem. They will be pleased with many different main courses, and many different bottles, but the best outcome requires that the meal and the wine match. This time, Alice chooses to buy a steak, rather than the salmon, which looks very good too. This is a public contribution to their activity; it goes on the record that Ann has bought the steak; and Bob knows it. It follows that Bob must buy red wine, rather than white. Again, this is because of the state of the world and Alice and Bob’s preferences about dinner, not because of a *signal* from Alice. Bob’s reason for buying red wine is that he knows Alice has bought steak, not that he recognizes Alice’s intention for what he should do. Bob in fact chooses a fine Bordeaux, so Alice and Bob solve their coordination problem.

The structure of this interaction highlights the difference between solving coordination problems and making meaning. Alice does not *mean* that Bob should buy red wine by her purchase of steak. She is only doing her part for dinner. Her purchase figures in the solution of their coordination problem and has real-world consequences, including updating the state of Alice and Bob’s joint activity. But Alice and Bob are not coordinating on what information should characterize the state of their joint activity, given what they have done: and only coordinating on adding information to the record underwrites attributions of meaning.

In fact, our utterances often contribute to realizing our practical projects, just like any other instrumental action. Thus, following Austin (1962), we distinguish the *illocutionary effects* our utterances conventionally engender from the *perlocutionary effects* they bring about through other causal mechanisms. “Saying something will often, or even normally, produce certain effects upon the feelings, thoughts, or actions of the audience, or of the speaker, or of other persons” Austin writes, “and it may be done with the design, intention, or purpose of producing them” (1962, p 101). Austin’s familiar examples include convincing, persuading, deterring, alarming, or surprising the audience. Such effects depend not only on the content that utterances contribute to the conversation, but also on the circumstances that obtain, on our preferences for the future, and on the flow of our ideas and emotions. We do not coordinate on these broader mechanisms. They are general and in no way conventional. Austin highlights this generality: “It is characteristic of perlocutionary acts that the response achieved, or the sequel, can be achieved additionally or entirely by non-locutionary means: thus intimidation may be achieved by waving a stick or pointing a gun” (1962, p 119).

We may coordinate on perlocutionary effects, of course, just as we coordinate on any other outcome we value. You may be working to convince me, and I may be working to

be convinced. But this is a change in me, occasioned by the make-up of my psychology; it is not *information* contributed to our interaction *through* our coordination. Your argument does not *mean* that you convince me, even if you do.

We can now see what's at stake in positing metaphorical meaning. A speaker successfully communicates *p* with an utterance only if he intentionally elicits an effect given by *p* in his audience by means of recognizing the speaker's intention to elicit this effect; or if the utterance contributes to coordinating with the audience to add *p* to their conversational record. If metaphors contribute content, then, this content must be identified either with what a speaker intends her audience to recognize her as trying to communicate; or as what must be added to the conversational record in consequence of coordination between her and her audience.

Metaphorical *meaning* cannot be merely content 'gotten across' nor can it be merely behaving recognizably as coordinating. It requires an audience to recognize the specific content a speaker wants to get across, and to use the signal of the metaphor as the basis for the uptake of that content. Since we deny this must happen in normal confrontations with metaphors, we therefore reject metaphorical meaning.⁵

We are *not* denying that metaphors can be used with the intention of drawing a hearer's attention to similarities, but this is *not* an intention to convey propositional content. We concede that a metaphor succeeds, if it does, from the cognitive effort an audience puts into exploring the similarities suggested by the metaphorical imagery, but we deny that this derives from recognizing a speaker's intention to convey propositional content.⁶ Taking this perspective allows us to afford a specific place for metaphor in the architecture of cooperative interaction *without* ascribing any meaning to a metaphor or its uses beyond the literal.

How are we to take stock of this proposal? We have thus far considered several cases whereby utterances bring about effects that are not part of what the utterances or their creators mean. None of these cases provides a close analogue to what happens in metaphor. The so-called secondary effects of utterances are a poor fit. When a speaker uses an English utterance, his audience learns that the utterance has been made; they learn that the speaker speaks English. Their inferences track what they know must be so when a speaker uses certain words. But appreciating a metaphor represents a different kind of inference. If we take away from (1) that love follows a trajectory from excitement, to misadventure, to torment, it's not because we think love *must* follow this trajectory for Groening to have formulated and produced his utterance.

Familiar cases of perlocutionary effects also provide a poor fit. When a speaker's utterance intimidates the audience, it's because its content triggers certain thoughts and emotions in the audience. The psychological dynamics run their course. The audience

⁵ Thanks to Zoltan Szabo for emphasizing the importance of this distinction to us.

⁶ To say metaphors make us see similarities is not to say metaphorical content is a simile, though we may be able to describe the effects of a metaphor with a simile. But of course this no more requires the metaphor to mean that one thing is like another than Herod's showing John heads requires us to have his act speaker mean that John is dead.

need not recognize the speaker's intention to intimidate – if the speaker even has such an intention – nor must the audience *participate* in any special way in calling up the thoughts and feelings involved, in bringing the effects about. By contrast, we cannot appreciate the suggestions of (1) – e.g. *that love begins its trajectory with excitement* – merely by submitting the proposition *that love is a snowmobile* to our ordinary faculty of inference, critical thinking, and deliberation. Love is not a snowmobile! To suggest otherwise is absurd at face value. Reasoning normally from (1) would just rediscover the absurdity. We must *recognize* that (1) is a metaphor, and *shape* our psychological response accordingly.

Other sorts of acts, however, upon reflection *do* offer closer parallels with metaphor. By drawing these parallels explicitly, we hope to clarify the status we are ascribing to metaphor.

Davidson likens speaking metaphorically to joking. Consider the conversation in (4), as reported by Lampert and Ervin-Tripp (2006, p 54).

4. *Friends talk about Diet Pepsi.*

- Vi: There is a problem with aspartame, but I can't remember what it is.
Joe: It affects your memory.
Vi: Hah! That's a good one.

As Vi and Joe spoke, Vi was drinking a Diet Pepsi. The drink is sweetened with aspartame. Vi takes Joe to be joking here. The sentence that he uses clearly *means* that aspartame affects memory, but his *point* in using it is *not* to contribute this information seriously to their interaction. In fact, despite Joe's utterance, it seems he makes no commitment whatsoever, one way or another, as to whether aspartame affects memory, or even whether it is associated with any other health problems.

What gives Joe's utterance its humor is the imagery behind it. The joke conjures up the imaginative possibility that Vi is already suffering from the ills of aspartame, but hasn't yet noticed the symptoms or their connection to the beverages she drinks. The effect is to ridicule whatever worries might lie behind Vi's broaching the topic of the problems of aspartame. Of course, in order to get Joe's joke, Vi must recognize that he is offering his utterance merely to develop the imagery behind the joke, not to contribute this information. She must participate in exploring the consequences of this imagery to discover that it portrays her as suffering from the ills of aspartame, and so, to discover that this joke offers Joe a way of chiding her for worrying about it. Thus, the joke requires special cognitive effort from Vi, prompted in part by her recognition of its status as a joke.

It seems clear that the effects that Joe achieves with (4) do not fit the account of speaker meaning from (2) or (3). Joe's goal here seems to be to ridicule Vi, not to assert that (he thinks) she is ridiculous. This outcome – creating a specific kind of interaction between people – is simply the wrong kind of effect for speaker meaning. Moreover, intention recognition plays a very limited role in the elicitation of this effect. It ends once Vi has

recognized Joe's utterance as a joke. From then on, the humor in the utterance is determined by the imaginative possibilities it portrays, not by Joe's intentions, Joe and Vi's mutual expectations, or the rules of language. These considerations seem so overpowering that we think few people would even be tempted to say that Joe's joke has a special, non-literal meaning, to the effect that Vi's concern with aspartame is laughable. If speakers mean anything by their jokes, they mean nothing other than what their words mean.

The situation is exactly parallel, we claim, with metaphor. Groening's metaphor conjures up the imaginative possibility that love is a snowmobile. To get its point, a reader must recognize Groening is offering his metaphor merely to develop the imagery behind it, and of course, not to contribute its literal meaning to the conversational record. On confronting the metaphor, an audience must participate in exploring the implications of its imagery. That effort may lead the audience to appreciate certain similarities between love and a tragic snowmobile ride: exhilaration, misadventure and torment. The goal of the utterance is for this appreciation to occur, not for specific information to be exchanged, and interlocutors do not coordinate on the information itself or derive it directly by intention recognition. Thus, even when a metaphor succeeds, it does not result in special, non-literal meaning. The parallel between metaphors and jokes, and the obvious appeal of the explanation as it applies to jokes, adds to the plausibility of our view of metaphor.

Analogously, consider hinting. Here is an instructive if not subtle example. A student brings a completed exam to the examiner's desk at the final for a small, friendly class. The examiner looks it over, finds an error and responds with (5).

5. You might want look at your answer to problem three again.

Under normal circumstances, an examiner would say (5) only because the student's answer to problem three is wrong. The student knows this. In fact, in this case, the examiner intends that the student revisit her answer, notice her error and correct it. The student knows this as well. But from this it does not follow that the examiner speaker *means* by his utterance of (5) that the student's answer to problem three is wrong. Indeed, if the examiner intends that the student should *discover* the error, then both the analyses in (2) and (3) predict that this is not a case of speaker meaning. The student finds that the answer is wrong by looking at her paper, not by recognizing the examiner's intentions.

This is a coordination problem. Both student and examiner prefer the student to fix her answer, and so, do better on the exam. It is likely the examiner has many ways to clue his student in, and the student has many strategies for pinpointing and fixing her error. But the relevant coordination, like Alice and Bob's over dinner, involves action in the world – in this case, particularly, the cognitive and problem-solving actions undertaken by the student.

In interpreting (5), the student must recognize the examiner's intention and act based on that recognition. Plausibly, this is an indirect instruction, and by (5), the examiner does

mean for the student to revisit problem three. Ultimately though, the student discovers the error on her own. So, the theory articulated in (2) or (3) predicts that the examiner does not mean with an utterance of (5) that problem three is wrong. He merely *hints* it is wrong.

Most hints are not as explicit as (5) in what they convey. The speaker prompts the hearer to reconsider some question. It really is up to the hearer what she learns from this inquiry. This gives most hints the open-ended texture found also in metaphor. If one has a specific message to impart, a hint is often a poor way to get it across. D. Stone, Patton and Heen (1999) make this point particularly effectively, using example (6), spoken by a wife to a husband who has taken up golfing every Saturday morning.

6. Is golf so important that you need to play it this often?

They write:

This is a classic example of a statement masquerading as a question. It's obvious that the meaning of the comment is conveyed in the subtext. What is less obvious is what the meaning is supposed to be. Your tone conveys anger or frustration. But it's not clear what is causing the anger or what your husband is supposed to do about it. Are you angry that your husband is engaged in a meaningless sport rather than community service or household chores? Are you angry that he's not taking you along? Are you angry that you aren't spending enough time together? How would he know? (D. Stone *et al*, 1999, p.192).

Ultimately, in hinting, a speaker does not achieve speaker meaning. A speaker invites his audience to participate in a certain kind of inference: thinking an issue through again, in light of new evidence that other views might be possible. This process may lead the audience to notice new things; it is the speaker's aim that such discoveries occur, but they are outside the realm of the speaker's control – and indeed that of the audience. The audience makes the discoveries that they are led to by the critical thinking they do. They do not make these discoveries merely in virtue of recognizing their interlocutors' intentions that they do so.⁷

⁷ Some of what Grice called conversational implicatures turn out to be more like hinting than speaker meaning. He introduces the term "implicature" neutrally:

A asks B how C is getting on in his job, and B replies, 'Oh quite well, I think; he likes his colleagues, and he hasn't been to prison yet.' At this point, A might well inquire what B was implying, what he was suggesting, or even what he meant by saying that C had not yet been to prison. The answer might be any one of such things as that C is the sort of person likely to yield to the temptation provided by the occupation, that C's colleagues are really very unpleasant and treacherous people, and so forth. It might, of course, be quite unnecessary for A to make such an inquiry of B, the answer to it being, in the context, clear in advance. It is clear that whatever B implied, suggested, meant in this example, is distinct from what B said, which was simply that C had not been to prison yet. I wish to introduce, as terms of art the verb 'implicate' and the related nouns 'implicature' (cf 'implying') and 'implicatum' (cf. 'what is implied'). The point of this maneuver is to avoid having, on each occasion, to choose between this or that member of the family of verbs for which "implicate" is to do general duty (1989, p. 24)

Grice has framed this discussion explicitly to avoid the question of whether implicatures are (speaker)

We draw the same conclusion about the metaphorist. While the point of Romeo's line 'Juliet is the sun' might be to prompt a reader to think of Juliet as radiant and centered, there should be no temptation to claim the metaphor has a "special" meaning to this effect. Similarly, it is a mistake to claim that when Auden writes, "The hourglass whispers to the lion's paw," what he thereby says has a "special (non-literal) meaning"—that time, like the steady flow of sands in an hourglass, is subtle and unrelenting. There is no such meaning. In short: the point a metaphorist uses a metaphor to make is not a property of the metaphor itself.

2 Criticisms of Metaphorical Meaning

This ends our outline of the positive proposal; before advancing it further, we need to see what's wrong with its alternatives. Every account of metaphorical meaning we are aware of has been criticized in one manner or another, but nothing like the criticism launched by Davidson. His central claim in "What Metaphors Mean" (1979) is that metaphors have only literal meaning. He does not question their potential effectiveness; only how these effects might be achieved. Davidson goes further than most critics in denying that there is a determinate propositional thought a speaker can *reasonably* be taken to mean by his use of a metaphor. A metaphor might "*make us* attend to some likeness, often a novel or surprising likeness, between two or more things;" it might make us "see one thing as another,"⁸ but this doesn't distinguish metaphor from a bump on the head or taking a drug: each can cause certain effects, like noticing surprising similarities. What he denies is that the metaphor, like the bump and the drug, *means* these effects.

For Davidson, what metaphors convey is largely *non-propositional*, unpredictable causal effects of their utterance-events understood exclusively according to their literal meanings. This doesn't mean that when Shakespeare wrote "Juliet is the sun," he had *nothing in mind*. Davidson is denying that whatever he had in mind must be part of a metaphorical interpretation his audience is expected to retrieve. He denies that, associated

meant. It's reasonable to conclude he thought many or even most implicatures were, but not all. Later, in considering irony, he appeals to considerations that seem to distance the interpretation of irony from the factors that go into his typical discussions of speaker meaning.

To be ironical is, among other things, to pretend (as the etymology suggests), and while one wants the pretense to be recognized as such, to announce it as a pretense would spoil the effect. What is possibly more important, it might well be essential to an element's having conventional significance that it could have been the case that some quite different element should have fulfilled the same semantic purpose; that if a contemptuous tone does conventionally signify in context that a remark is to be taken in reverse, then it might have been, for example, that a querulous tone should have been used (instead) for the same purpose. But the connection of irony with the expression of feeling seems to preclude this; if speaking ironically has to be, or at least appear to be, the expression of a certain feeling or attitude, then a tone suitable to such a feeling or attitude seems to be mandatory, at least for the unsophisticated examples (1989, p. 54)

Thomason (1990) says "most implicatures are meant". He goes on to explicitly consider the relationship between implicatures and Grice's definitions, to show that many implicatures come out as meant. We endorse this understanding of implicatures and emphasize that our arguments do not count against the meaningfulness of most implicatures, particularly Grice's Group A implicatures, where no maxim is flouted. This category includes most bread-and-butter examples of relevance and quantity implicatures: those associated with inferring coherent interpretations in discourse.

⁸ Davidson (1978), pp. 31, 45; our emphasis.

with his (or any other metaphor), there is a content every successful interpreter must grasp (Davidson, 1979, p 425). His conclusion is based on several arguments; given the nature of speaker meaning and our positive proposal for what metaphors do, it will be useful to rehearse some of these criticisms here.

2.1 Ineffability of Metaphorical Content and Rejoinder

Davidson writes:

If a metaphor has a special cognitive content, why should it be so difficult or impossible to set it out? If, as Owen Barfield claims, a metaphor “says one thing and means another,” why should it be that when we try to get explicit about what it means, the effect is so much weaker—“put it that way,” Barfield says, “and nearly all the tarning, and with it half the poetry, is lost.” 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? (pp. 423–424)

Davidson answers his questions:

The reason is ...that we imagine there is a content to be captured when all the while we are in fact focusing on what the metaphor makes us notice. If what the metaphor makes us notice were finite in scope and propositional in nature, this would not in itself make trouble; we would simply project the content the metaphor brought to mind onto the metaphor. But in fact there is no limit to what a metaphor calls to our attention, and much of what we are caused to notice is not propositional in character...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 (pp. 425)

If a metaphor has a speaker meaning, it should be possible to express it literally. Otherwise, we are left with the puzzle of how a speaker could formulate, and an addressee could recognize, an intention to convey a specific proposition, without any independent way of specifying exactly what proposition that was. Such puzzles simply do not arise if we attribute other kinds of intentions to a metaphorist, such as an intention to get an audience to notice similarities.

One reason why it’s so difficult to paraphrase the metaphor into literal language is because it can legitimately ‘say’ different things to different hearers. In this regard, metaphorical interpretation, for lack of a better phrase, is a creative open-ended affair. When Matt Groening says “Love is a snowmobile,” it occurred to one of us that the way cold air takes your breath away is a lot like the shock of being around somebody you have a crush on. That never occurred to the other one of us. He was instead thinking of the speed, the smooth ride, the unfamiliar environment. Still, the first impression is just as valid as an insight the metaphor delivers. But Matt Groening’s goal was not to get us

to grasp a deep subtle truth, but rather to see something in a certain way, and this is *not* something speaker meant. As Davidson explains:

...if I show you Wittgenstein's duck-rabbit, and say, 'It's a duck', then with luck you see it as a duck; if I say, 'It's a rabbit', you see it as a rabbit. But no proposition expresses what I have led you to see. (pp. 425)

The difficulty in devising a literal paraphrase of what a hearer has been led to see is *not* that the latter is especially subtle or profound, but rather, that seeing a figure in this (or any other) way is *not* amenable to literal paraphrase (though of course it can be described).

Some authors, Moran (1996), deny that all "cognitive contents" can be construed propositionally. Metaphor might represent the world non-sententially, and so, are no more amenable to literal sentential articulation than a painting of a landscape is; and so, to arrive at an interpretation of a metaphor we are not required to grasp a meaning beyond the literal. We needn't grasp a proposition to whatever effect the metaphor prompts us to see.

Even if Moran is correct about the ontology of content, his proposal concedes that a speaker cannot successfully use metaphors to intentionally elicit an effect given by a propositional content *p* in his audience by means of their recognizing this intention to elicit this effect; nor does it permit the metaphorical use to contribute to coordinating with an audience to add the propositional content *p* to their conversational record. That is, Moran (and others) conceding that however metaphorical use contributes cognitive content to a discourse it is not through speaker meaning. Non-sentential representation – content not grounded in a proposition – is not of the right shape for this.

2.2 Metaphors and Similes

Davidson's next argument compares metaphors to similes.

...how is it that a simile gets along without a special intermediate meaning? In general, critics do not suggest that a simile says one thing and means another—they do not suppose that it means anything but what lies on the surface of the words. It may make us think deep thoughts, just as a metaphor does; how come, then, no one appeals to the "special cognitive content" of the simile? (Davidson, 260–261)

Just as we don't look for hidden meaning with similes, why should we with metaphors? There seems, however, to be a difference: Romeo did *not* mean Juliet is identical to a celestial body. So, this argument goes, he must have *meant* something else. And so, what he intended to communicate must be different from what his words literally express. And it's this "something else" that is its metaphorical meaning. So, even if similes can "work" without "special cognitive content," the same conclusion should not be drawn about metaphors.

Davidson, in turn, wonders why we must assume there is something a speaker who uses a simile *means*. Why isn't its purpose, Davidson asks, just to get a hearer to *notice* similarities; not to mean something? The metaphorist's intention needn't be to communicate a proposition affirming noticed similarities, but rather for his hearer to understand the proposition literally expressed; and then, through that understanding perhaps be *prompted* (contingent on his background information) to look for certain similarities. Davidson says, "a simile tells us, in part, what a metaphor merely *nudges* us into noting"; metaphors "intimate what the simile declares" (Davidson, p. 253).

When Romeo says, speaking metaphorically, "Juliet is the sun," Shakespeare expects his audience to grasp the literal meaning, and thereby, be prompted to notice similarities between Juliet and the sun. As with simile, no "proposition" is meant. It might well be that the speaker herself believes various propositions; she may even convey them with her utterance. But, so Davidson claims, if she does, it is only as a "by-product" of getting her audience to *notice* similarities. In short, there is no basis for claiming propositions are speaker meant by her utterance.

Davidson writes, "Metaphor and simile are merely two among endless devices that serve to alert us to aspects of the world by inviting us to make comparisons." The difference between simile and metaphor – as opposed to anything else that gets us thinking about similarities – according to Davidson (p. 256) – is that is that whereas similes literally tell us about resemblances, and other forms of expression can direct our attention to similarities, metaphors *bully* us into making comparisons.

2.3 Our Relationship to Davidsonian Views

Davidson's critical discussion is telling. But his positive view is merely a sketch, presented through analogies that are provocative but, we find, misleading. To liken metaphor to a drug or a bump on the head, for example, suggests a model where metaphor falls in with perlocutionary effects like intimidation. And as we have seen, this account of the relationship between pragmatics and psychology in metaphor is unsatisfying, because it doesn't explain the audience's need to recognize the metaphor and participate in exploring it.

Other philosophers also eschew metaphorical meaning, in ways consistent with Davidson's criticisms. Rorty (1987), for example, describes metaphors as "unfamiliar noises", and emphasizes that an audience, confronted with a new metaphor, undertakes an active and creative effort of interpretation. This process distinguishes metaphor from meaning; it involves no conventions, no signaling. However, Rorty locates this interpretive effort as part of a continuum of more or less constrained problems of interpretation. His continuum runs over many of the distinctions we have charted, and so, in some respects, it too obscures the relationships between pragmatics and psychology in metaphor. The metaphorist intends his audience to shape their interpretation in distinctive ways appropriate to metaphor, and they can appreciate the point of the metaphor only by following through. Rorty locates birdsong in this continuum, for example. Hearing an unfamiliar howl in the jungle, we actively explore interpretations that explain what we must be hearing. But this is to liken metaphorical interpretation to

recognizing Stalnaker's secondary effects: what must be the case for an utterance even to be made. This is qualitatively different, we have seen, from metaphor.

In short, giving a specific locus for metaphor in pragmatics and psychology, not just in an undifferentiated wastebasket of "use", gives us a conceptual framework, empirical analogues and technical resources for thinking about metaphor in precise new ways. Our view thus allows us to endorse many of the criticisms Davidson marshals against metaphorical meaning, while responding in new ways to the objections others level against him.

3 Metaphor without Meaning

There are positive arguments in favor of metaphorical meaning, and critics have attacked the Davidsonian position on various grounds.⁹ Most of them we believe miss their target but one that most critics urge is that the account conflicts with the intuitions about how metaphors are organized, understood and used (Bergmann 1982; Hills 1997; Stern 2000). We respond to these criticisms here.

3.1 Metaphor and semantics

First, there seem to be cases where metaphors interact by compositional semantics with the interpretation of other elements in a sentence. How can metaphors do such a thing unless they have meanings? We think the data comes in two families which require different responses. One family is illustrated by the famous quote by Donne in (7).

7. No man is an island entire of itself; every man
is a piece of the continent, a part of the main.

With (7), of course, Donne compares the social relationships among people to the geographical relationships among places. The metaphor seems to give a new interpretation to the word "island", associating it with the property of being so socially isolated as to be devoid of connections to other people. The first clause of (7) denies that any man has this property. How are we to square this interpretive interaction with our claim that there is no metaphorical meaning, no new sense of "island" at play in (7)?

We deny that the metaphorical thinking involved in understanding (7) takes scope with respect to the semantic constructs of Donne's description. Instead, we suggest, the reader's engagement with (7) involves thinking of people as places once and for all, and interpreting Donne's words, with their literal meanings, in the landscape that results. There are no islands, the reader finds, in this fantasy, only a patchwork of regions, identified as people in our imaginative correspondence, forming a single vast continent. This way of looking at metaphorical interpretation makes clear that the word "island" is not the unit of Donne's metaphor. The entire discourse is. Accordingly, this way of looking at metaphor dispenses with any metaphorical meaning for Donne's words, as we say it must.

⁹ See Moran (1989 and 1996) for these and other objections.

Though example (7) involves negation, it's easy to see how this explanation would extend to other operators that take scope within a clause, including tense, modality, aspect, quantification, and others. Since we can imagine, in any of these cases, what the world would have to be like for the clause to be true, we can actively undertake to imagine things *that way*, against the background of an imaginative correspondence initiated by metaphorical thinking.

The second family of potential counter-example is found in attitude reports such as (8)

8. Does Chris really think no man is an island? I don't.

In (8), the speaker is not just asking if Chris subscribes to the literal proposition that no men are islands. Rather, the speaker seems to be asking if his addressee accepts the *point* of the metaphor, that people are all interconnected by their social relationships. This requires the point of the metaphor to be part of the content of the thought. With metaphorical meaning, we have a ready explanation: take "no man is an island" semantically to contribute the proposition that people are all interconnected by their social relationships. But how can the semantics derive this requirement on our view, where the point of the metaphor is not a matter of semantics?

We propose instead to refine our understanding of the meaning of 'think' and other propositional attitude words. With (8), it seems that the addressee cannot truthfully answer "yes" just because Chris agrees that people are all interconnected by their social relationships. Answering "yes" also requires that Chris accepts the metaphor as apt, and moreover that Chris is drawn from there by metaphorical thinking to appreciate that people are all interconnected by social relationships. The metaphor must be active in Chris's thought, and so it must somehow also be somehow active in the truth conditions of (8).

How can we capture the active role of metaphor? We have argued that our insights in metaphorical thinking are prompted just by the literal meanings of utterances. Analogously, we will argue that statements such as (8) also *report* metaphorical thinking, and the insights that follow from it, through constructions that appeal just to the literal meanings of the utterances with which the thoughts are framed. In other words, a sentence of the form "A thinks that *p*", where *p* is normally understood metaphorically, is true if and only if *A* endorses the similarities *A* discovers upon entertaining the metaphor *p*. Finally, then, though (8) may tell us something about the role of indirection and metaphor in the meaning of the word "think", it tells us nothing about indirection and metaphor in the meaning of "no man is an island".

3.2 Metaphor and pragmatics

Even if we are right, and compositional semantics does not require metaphorical meaning, others have argued that *pragmatics* requires metaphorical meanings. We can give instructions with metaphors, for example. Swimming instructors are fond of using imagery like (9) to teach the butterfly, the strenuous but beautiful stroke where swimmers

must lift their shoulders enough on the recovery to swing both arms at once forward over the water.

9. With each stroke, let a wave flow along the length of your body, from your head down your back and through your legs.

The instruction does not advise the swimmer literally to bob along with the surface of the water in the pool, but invites the swimmer to guide their action by the image of a figurative wave: the sinuous flow the swimmer creates in her own body. How does an addressee know what to do, on hearing (9), without extracting an action description that counts as the meaning of the coach's metaphor? Indeed, how else does the addressee even realize that a metaphorical instruction such as this gives a real command?

Such objections fail to confront the richness with which our imaginative efforts can guide our real-world actions. Problematically, in swimming the butterfly, it is no use trying to lift the shoulders outright. The hips fall, the body drags, the arms pull less effectively, and, for all that, the shoulders remain too low in the water, as they rise only as much as the swimmer's natural buoyancy permits. The image of (9) reorients the butterfly recovery as a dynamic process. It suggests to the swimmer that she can lift the shoulders for a brief moment in the cycle of the stroke by harnessing the strength of the abdomen and the propulsion of the feet, then goes on to hint how she can maintain an efficient and fluid rhythm in the stroke by accepting and relaxing into the downward reaction in the head and shoulders that inevitably follows. She can use the image of a wave to guide the orchestration of her movements. Such an imaginative effort can result in swimming the butterfly, even when the swimmer herself cannot conceive of a literal description of the movements that constitute that action. When a swimmer knows her coach demands both her imaginative and her real-world effort, there seems no difficulty for her to play along with (9) as an instruction, and swim in response.

More generally, we think, objections to a Davidsonian account of metaphor on pragmatic grounds miss the extent to which non-serious, imaginative and indirect language often gathers a life of its own, in the collaborative action of interlocutors, across extended patterns of interaction. Humor shows this most clearly. Examples (10) and (11) are edited transcripts from *The Colbert Report*, omitting disfluencies and side conversations that are irrelevant to our point.

10. *Stephen Colbert and Penn Gillette, on cold hard reality*

- C: But why would you break the hearts of children who want to believe there is magic in the world?
- G: [to the camera] There is no *Santa Claus* [bleeped]. The tooth fairy is *your mother and your father* [bleeped]. There is no *Easter bunny* [bleeped], and, no God. Sorry.
- C: Do we have a puppy for Mr Gillette to punch, before we leave here today?
- G: I don't like to touch them. I like to hit them with a hammer.

In (10), Stephen Colbert teases magician Penn Gillette, of the performance team Penn and Teller, about his strategy of giving away magicians' secrets. Colbert offers humorous

imagery to describe the pain Gillette might inflict by revealing the information he does. Gillette “plays along” by jokingly offering his own mock demonstrations of cruelty. With his final comment, Colbert compares telling children the truth to punching puppies. Gillette’s response endorses this cruel image, and takes it further. He gets no visceral satisfaction, he jokes: the cruelty itself sustains him.

In (11), Colbert’s banter with documentary filmmaker Morgan Spurlock hints at the subversiveness of imagery that allows people to understand one another better.

11. *Stephen Colbert and Morgan Spurlock, on empathy*

- S: We kind of live in a world where we don't really know what other people's lives are like, and by doing this [Spurlock's TV show, 30 days] you kind of immerse yourself in an environment where you learn a little bit along the way.
- C: But don't we live our own lives so we don't have to know what other people's lives are like? I mean, that's why I've got power windows, so I can roll up the window when I go through the neighborhoods I don't live in.
- S: I think we need one in between the seats too, so you can't see the person next to you.
- C: That'd be fantastic. Maybe our own oxygen system.

In response to Spurlock’s description of his new show, Colbert takes up his character of conservative talk-show host to playful extremes with his imagery of disengagement. Spurlock follows along, pushing Colbert to a *reductio* which pictures everyone living alone in a bubble.

We think such examples are commonplace, not just for jokes, but for metaphor, hints, and other forms of imaginative and indirect discourse. We emphasize both the ordinariness and the playfulness of interlocutors’ language. Gillette, in the final utterance of (10), for example, presents an utterance that purports, in the ordinary way, to contest a presupposition of Colbert’s prior utterance: that Gillette expects to punch puppies. Gillette avers that he will smash, but not punch. This is an entertaining development of the interlocutors’ theme – in fact, their metaphor – of truth as cruelty. It should not be taken to indicate that Colbert really did presuppose that Gillette would expect to punch a puppy, nor that Gillette really does like to hit puppies with hammers, nor that Colbert and Gillette have a real dispute over what Gillette likes to do with puppies.

Likewise, Colbert and Spurlock in (11) present utterances that purport, in the ordinary way, to develop a solution to a collaborative planning problem: how best to avoid contact with people unlike ourselves. But Spurlock’s *point*, clearly – and perhaps Colbert’s as well – is not to solve this problem, but to ridicule it. In no way is it a goal worth solving.

In (10) and (11), interlocutors use the language of debate, but in fact what they are doing is developing imagery and cementing a shared understanding of their situation, actions and values. They do not contest propositions. In much the same way, we expect, interlocutors use their metaphorical discourse not to assert and deny propositions but to develop imagery and to pursue a shared understanding. Such practices can account for

our interactions in using metaphor, without appealing to metaphorical meaning or metaphorical truth. Indeed, they are all the stronger for dispensing with such notions.

3.3 *Metaphor and understanding*

Whether or not we need metaphorical meaning for semantics or pragmatics, some have argued we still need metaphorical meaning to describe the knowledge and processes involved in *understanding* a metaphor. Metaphorical interpretation does rely on background knowledge. No one would interpret an utterance of “Man is a wolf” as *men are vegetarians*. Anyone interpreting the metaphor as such would either, as Black says, *not be a member of our speech community* or fall short of our expectations of the common knowledge of those members; he would “produce an effect of paradox and provoke a demand of justification” (Black 1955).¹⁰ Indeed, there is something culturally-defined in the “system of associated commonplaces” which seems to guide listeners’ imaginative efforts in metaphor. To say of John “He is a gorilla” is to invite comparisons with the gorillas of myth and stereotype, and not the gorillas of nature.

We accept such observations, but they do distinguish sharply between the mechanisms of meaning and other tools we have for sharing information with one another. Returning to Herod, Salome, and the head of John the Baptist, for Salome to understand from Herod’s demonstration that John is dead, she must draw on background knowledge: severed heads are doomed to death. To fail to draw this conclusion would fall short of our expectations of the common knowledge of members of our speech community. Nevertheless, in light of the distinctions elaborated in §1, this is not to argue that Herod *means* by showing John’s head to Salome *that John is dead*. Black’s observation about wolves is thus no argument for metaphorical meaning either.

Commonplaces of the imagination, of course, are just as effective in prompting reactions to real-world stimuli as is knowledge of actual facts. A gun intimidates, in Austin’s example; it is genuinely dangerous. But threatening, say, to feed somebody grubs or insects might be perfectly scary or revolting – though they are perfectly safe as food, even a delicacy in certain parts of the world. In our way of life, we imagine them as foul, diseased, poisonous. To appreciate the power of the imagination in such cases, and to note that the imagination lives a life of its own, divorced from what we might know to be the case, is not to locate the force of the imagination in convention or meaning. It is no mere convention that offering the grubs revolts you. But we can say the same of the power of the imagined gorilla, in Black’s example. It is no argument for metaphorical meaning.

We often understand and at times misunderstand metaphors. Critics disagree about which interpretation is more accurate. The coherence of their disagreement suggests a “cognitive content” an audience might understand or misunderstand.

But how does such data support the existence of metaphorical meaning? Why isn’t what’s understood or misunderstood the *point* of the metaphor—that is, why isn’t it what

¹⁰ See Hills (1997) and Bezuidenhout (2001) for further discussion of this point.

the metaphorist wants the interpreter to *notice*, to *see* without communicating it to him. The sense in which metaphors are “cognitive” is that they can be understood or misunderstood; but this phenomenon is only “derivative,” as it were, they are misunderstood only in the sense that their point is missed.

Again, it’s certainly possible to mistake someone’s hint, or to fail to get their humor. Nobody would take such cases of misunderstanding as evidence that jokes, for example, carry indirect meanings. We insist on the parallel with metaphors.

3.4 Dead metaphors

Some authors have thought that dead metaphors pose a problem for any theory that denies metaphorical meaning. Everyone agrees that when a metaphor dies, it acquires a new literal meaning; and in normal circumstances this literal meaning is what a speaker means when he uses the dead metaphor. It has seemed plausible to many to suppose that this newly acquired meaning must have been *active* all along while the metaphor was alive. Without making some such supposition, the argument goes, how could one possibly account for the cognitive content’s current association with the dead metaphor? But then, isn’t it plausible to suppose the newly acquired literal meaning *is* the “cognitive content” of the once living metaphor that Davidson denies exists?

Why is this supposed to be a problem for Davidson? Suppose that when a metaphor dies, it thereby acquires a second literal meaning. But even if dead metaphors—when living—were used to introduce content, how does help the advocate of the argument from dead metaphors unless he can establish that the new literal meaning is what we speaker meant by uses of the then alive metaphor?

Davidson, in turn, asks us to look at those countless metaphors that are not destined to die. Why don’t they die? His response is that such metaphors do not die because they cannot; and they cannot because they are *not* used to speaker mean anything propositional. They are constitutionally incapable of dying. Thus, while some metaphors (those destined to die) might well be used to introduce some content, the same cannot be claimed for all of them.

Here’s Davidson’s own response:

If metaphor involved a second meaning, as ambiguity does, we might expect to be able to specify the special meaning of a word in a metaphorical setting by waiting until the metaphor dies. The figurative meaning of the living metaphor should be immortalized in the literal meaning of the dead. But although some philosophers have suggested this idea, it seems plainly wrong. “He was burned up” is genuinely ambiguous (since it may be true in one sense and false in another), but although the slangish idiom is no doubt the corpse of a metaphor, “He was burned up” now suggests no more than that he was very angry. When the metaphor was active, we would have pictured fire in the eyes or smoke coming out of the ears. (p. 419)

If metaphors involved “second meanings,” these would be the literal meanings acquired by metaphors upon death. But they are not, and so, metaphors lack them. One might challenge his premises by claiming only the cognitive aspect of a living metaphor emerges as the literal meaning of the dead metaphor or challenge his claim that the literal meaning of the dead metaphor *is* the “second meaning” of the living one. What distinguishes the living from the dead metaphor is that the former has non-semantic (perhaps imagistic) components not possessed by the latter. There is an important difference between the literal meaning a metaphor acquires upon its death and the “intimation” of a living metaphor. *But* this hardly challenges Davidson; it is *his* point! The “intimation” of a living metaphor appears different in kind from the newly acquired speaker meaning of a dead metaphor. To claim living metaphors have special “cognitive contents” distinct from their literal meanings is to suggest what is false: namely, that literal meanings and metaphorical “intimations” are alike.

4 Conclusion

In this paper, we have distinguished metaphorical thinking – developing imagery, seeing one thing as another, noticing similarities – from merely grasping a proposition, namely, the one that is speaker meant, brought about through an intention to present information through coordination or intention recognition. If we are right, then researchers who try to assimilate metaphor to other mechanisms for meaning making in language are on the wrong track. Relevance theorists (Carston, 1991; Bezuidenhout, 2001, Recanati, 2001), contextualists (Stern, 2000), and traditional pragmatic theorists (Grice, 1975, Searle 1979), each in their own way, develop an account of metaphor in familiar terms. To the extent that metaphor is a distinct process with a distinct character and distinct consequences, as we have argued, the uniformity of these accounts is a weakness, not a virtue.

Other researchers, however, give accounts of metaphorical meaning that do depend on a distinctive kind of metaphorical thinking. Hills (1997), for example, following Walton (1993), sees metaphorical thinking as prop-oriented make-believe. On Hill’s account, the resemblances we appreciate in metaphor help us to develop an imagined world where one thing takes the place of another. Camp (2003), meanwhile, describes metaphorical thinking as a kind of analogical reasoning. For Camp, by finding metaphorical parallels, we can focus on salient properties of familiar objects and use them to attend anew to the counterparts we discover for them in a new domain.

In this paper, we have said nothing to undermine such views of metaphorical thinking. We think addressees *might* recognize the point of a metaphor in the ways that Hills or Camp describe. We counter only the arguments Hills and Camp advance that the point so recognized is the *meaning* of the metaphor.

Why should we care if a point so recognized is a meaning? What’s in a name? Both meaning and metaphor sit at the center of complex webs of phenomena, principles and puzzles. There is much to gain from keeping the issues separate. In the realm of meaning, the hardest problems come in regimenting our knowledge of language, and thus coming to understand how that knowledge is acquired, structured, and used in speakers’ faculty

of language. Our intuitions about literal meaning place heavy constraints on key semantic notions: truth, reference, context and logical form, among others. Progress, we have argued (Cappelen and Lepore, 2004), depends on developing and applying clear standards that demarcate semantic phenomena narrowly. If we can locate metaphor elsewhere, it is good news for meaning.

The study of metaphor likewise benefits from putting issues of meaning aside. Philosophers have always wanted to know, not just what metaphor is, but what we do when we appreciate or craft a metaphor. It prejudices such questions to class metaphor as meaning: it suggests that we produce and understand metaphors in more or less the ordinary way we produce and understand all other language. Distinguishing the point of metaphor from meaning lets us settle these questions on their merits.

Since Aristotle's invocation of the special genius of the metaphorist, critics have worked to pin down the ingredients that make metaphor special. Rarely do they invoke the "prosaic" ideas of coordination, convention, or intention recognition. More often, they locate the power of metaphor in the real relationships among things that only metaphorical thinking can bring us to notice. Barfield (1928/1973) quotes Bacon on metaphor:

Neither are these only similitudes, as men of narrow observation may conceive them to be, but the same footsteps of nature, treading or printing upon several subjects or matters.

Barfield goes on:

This is the answer. It is these 'footsteps of nature' whose noise we hear alike in primitive language and in the finest metaphors of poets. Men do not *invent* those mysterious relations between separate external objects, and between objects and feelings or ideas, which it is the function of poetry to reveal. These relations exist independently, not indeed of Thought, but of any individual thinker. And according to whether the footsteps are echoed in primitive language or, later on, in the made metaphors of poets, we hear them after a different fashion and for different reasons (Barfield, 1928/1973 p 86).

Meaning, certainly, does not work this way. Metaphor might. Divorcing metaphor from meaning, as we advocate here, opens the door to honor, to refine and to defend such insights, in the terminology of psychology, action and information. Where it belongs.

References

- Owen Barfield. 1928/1973. *Poetic Diction: A Study in Meaning*. First published by Faber and Gwyer, 1928. Reference to Wesleyan edition, 1973.
- J. L. Austin. 1962. *How to Do Things with Words*. Harvard University Press.
- Merrie Bergmann. 1982. "Metaphorical Assertions." *Philosophical Review* 41:229–245.
- Anne Bezuidenhout. 2001. "Metaphor and what is said: a defense of a direct expression view of metaphor." *Midwest Studies in Philosophy* 25, 156-186.
- Max Black. 1955. "Metaphor." *Proceedings of the Aristotelian Society* 55:273-294.
- Elisabeth Camp. 2003. *Saying and Seeing-as: The Linguistic Uses and Cognitive Effects of Metaphor*. PhD Dissertation, UC Berkeley.
- Herman Cappelen and Ernest Lepore. 2004. *Insensitive Semantics: A Defense of Semantic Minimalism and Speech Act Pluralism*. Blackwell.
- Robyn Carston. 1991. Implicature, explicature, and truth-theoretic semantics. In S. Davis (ed.), *Pragmatics: A Reader*. Oxford University Press.
- Donald Davidson. 1979. "What Metaphors Mean." in *On Metaphor*, ed. Sheldon Sacks. Chicago: University of Chicago Press
- David Devault, 2008, *Contribution tracking: Participating in Task-Oriented Dialogue under Uncertainty*. PhD Dissertation, Rutgers.
- H. P. Grice. 1957, "Meaning" *The Philosophical Review* 64: 377-388.
- _____, 1969, "Utterer's Meaning and Intentions," *Philosophical Review* 78: 147-77.
- _____, 1975. "Logic and Conversation," *Syntax and Semantics, volume 3*, P. Cole and J. Morgan, eds. New York: Academic Press;
- _____, 1989, *Studies in the Way of Words*, Harvard University Press.
- David Hills. 1997. Aptness and truth in verbal metaphor. *Philosophical Topics* 25:1, 117-153.
- Martin D. Lampert and Susan M. Ervin-Tripp. 2006. "Risky laughter: Teasing and self-directed joking among male and female friends." *Journal of Pragmatics* 38:51-72.
- Ernest Lepore and Matthew Stone. 2006. "Logic and Semantic Analysis," *Handbook of the Philosophy of Science*. Volume 5: Philosophy of Logic (ed) Dale Jacquette, Elsevier, 2006, pp. 145-176.
- David Lewis. 1969. *Convention: A Philosophical Study*. Harvard University Press.
- _____. 1979. "Scorekeeping in a Language Game", *Journal of Philosophical Logic* 8 pp 339-59
- Francois Recanati. 2001. "Literal/nonliteral." *Midwest Studies in Philosophy* 25, 264-274.
- Richard Rorty. 1987. "Unfamiliar noises I: Hesse and Davidson on metaphor," *Proceedings of the Aristotelian Society* (supp. volume 61), 283-296.
- John Searle. 1979. "Metaphor." In *Expression and Meaning: Studies in the Theory of Speech Acts*. Cambridge University Press.
- Robert Stalnaker. 1978. "Assertion", in P. Cole (ed), *Syntax and Semantics 9*, New York: New York Academic Press, 315–332.
- _____, 1998. "On the representation of context". *Journal of Logic, Language and Information* 7:3-18
- Mark Steedman and Matthew Stone. 2006. "Is semantics computational?", *Theoretical Linguistics* 32:73-89.
- Josef Stern. 2000. *Metaphor in Context*. MIT Press.

- Douglas Stone, Bruce Patton and Shiela Heen. 1999. *Difficult Conversations: How to Discuss What Matters Most*. Penguin.
- Richmond H. Thomason. 1990. "Accommodation, meaning, and implicature: Interdisciplinary foundations for pragmatics". In *Intentions in Communication*, ed. P. R. Cohen, J. L. Morgan & M. Pollack, pp. 325–63. Cambridge, MA: MIT Press.
- Kendall Walton. 1993. Metaphor and prop-oriented make-believe. *The European Journal of Philosophy* 1:39-57.