

Wittgensteinian Semantics

Author(s): Michael Pelczar

Source: *Noûs*, Vol. 34, No. 4 (Dec., 2000), pp. 483-516

Published by: [Wiley](#)

Stable URL: <http://www.jstor.org/stable/2671879>

Accessed: 08-10-2015 03:23 UTC

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



Wiley is collaborating with JSTOR to digitize, preserve and extend access to *Noûs*.

<http://www.jstor.org>

Wittgensteinian Semantics

MICHAEL PELCZAR
University of Virginia

Abstract

Wittgenstein emphasizes two points concerning his notion of family resemblance. One is that the use of a family resemblance expression resists characterization by certain kinds of rules; the other is that due to the prevalence of family resemblance in the philosophical lexicon, philosophical inquiry must in many cases proceed differently from how it traditionally has. This paper develops an interpretation of family resemblance that seeks to do justice to these claims. I argue that what is characteristic about family resemblance expressions is not that they exhibit a basic semantic feature unique to themselves, but that they combine a number of semantic properties that happen not to be coinstantiated elsewhere. These features include (1) content variability (also a property of ambiguous expressions, polysemes, and standard indexicals), (2) a feature I call “topicality” (which is also a characteristic of polysemes), and (3) “semantic openness” (a feature of many ordinary indexicals). The notions of topicality and semantic openness are explained, and certain terms of natural language are shown to be family resemblance expressions. I conclude by indicating some of the potential philosophical ramifications of these results.

1. Introduction

In a series of passages now nearly as famous as Plato’s allegory of the Cave, Wittgenstein introduces an idea that lies at the center of much of his later work: the idea of “family resemblance.”¹ These passages are widely regarded as among the clearest and best understood of the *Investigations*;² yet, existing interpretations of them leave some hard questions unanswered. My goal here is to address these questions head on. If successful, this effort promises to yield a better understanding not only of Wittgenstein’s later work, but also of the semantics of natural language.

Under the standard interpretation, Wittgenstein’s claim in §§66–68 is that certain kinds are such that the conditions necessary and sufficient for member-

ship in them can vary from one member to the next.³ According to this interpretation, Wittgenstein claims that there is no set of conditions necessary and jointly sufficient for being a member of any such kind—e.g., for being a game, or a number. In what has since become a popular way of presenting the standard interpretation, Renford Bambrough sets it forth as follows:⁴

We may classify a set of objects by reference to the presence or absence of features *ABCDE*. It may well happen that five objects *edcba* are such that each of them has four of these properties and lacks the fifth, and that the missing feature is different in each of the five cases. A simple diagram will illustrate this situation:

<i>e</i>	<i>d</i>	<i>c</i>	<i>b</i>	<i>a</i>
<i>ABCD</i>	<i>ABCE</i>	<i>ABDE</i>	<i>ACDE</i>	<i>BCDE</i>

Here we can already see how natural and proper it might be to apply the same word to a number of objects between which there is no common feature. And if we confine our attention to any arbitrarily selected four of these objects, say *edca*, then although they all *happen* to have *B* in common, it is clear that it is not in virtue of the presence of *B* that they are all rightly called by the same name.

The standard interpretation appears to face an obvious problem, one in fact that Wittgenstein himself raises, when he has his interlocutor object: “All right: the concept of number is defined for you as the logical sum [i.e., disjunction] of these individual interrelated concepts.” In terms of Bambrough’s presentation, the objection is that Wittgenstein’s observations fail to support the claim that there is no condition (or, set of conditions) necessary and sufficient for being a game, since for all that has been shown, the following condition counts as such:⁵

$$ABCD \vee ABCE \vee ABDE \vee ACDE \vee BCDE$$

The objection seems to make a lot of sense. Why can’t we just break the problem down into cases, giving necessary and sufficient conditions for chess, similar conditions for baseball, Poker, tic-tac-toe, and so on, and then define “game” disjunctively in terms of these various sets of necessary and sufficient conditions? Is it because the concepts of “being a game of chess,” “being a game of baseball,” etc. are vague? But Wittgenstein apparently does not take himself simply to be reintroducing the concept of vagueness under a new name; there are many expressions (such as “bald”) that, while vague, do not seem to be the sorts of terms that Wittgenstein has in mind when speaking of family resemblance.⁶ But then why can’t we say that something is a game just in case it satisfies the condition consisting of the logical disjunction of all those sets of necessary and sufficient conditions?

Here is Wittgenstein’s rather sketchy response to this question:

—It need not be so. For I *can* give the concept “number” rigid limits in this way, that is, use the word “number” for a rigidly limited concept, but I can also use it so that the extension of the concept is *not* closed by a frontier. And this is how we do use the word “game.” For how is the concept of a game bounded? What still counts as a game and what no longer does? Can you give the boundary? No. You can *draw* one; for none has so far been drawn. (But that never troubled you before when you used the word “game.”) (§68)

The standard interpretation sometimes tries to avail itself of this reply by extending Bambrough’s device as follows: in constructing the series of conditions, one can never stop at a certain point and declare that no further conditions are to be listed:

$ABCD \vee ABCE \vee ABDE \vee ACDE \vee BCDE \dots$

But here the schematic character of Bambrough’s exposition forcibly makes itself felt. What we need is a clear way of understanding the schema—of understanding what it means to say that one *can’t* (legitimately?) ever stop adding to the list. Does the ellipsis represent a special metaphysical property of family resemblance kinds? Or does it correspond to a feature of the use of a family resemblance term? And if the latter, what feature?

Commentators have supplied a variety of suggestions as to how the ellipsis should be understood; unfortunately, none of these proves to be satisfactory. According to Baker and Hacker, for example, Wittgenstein’s reply amounts to this: the extension of (e.g.) the word “number” is “open, not sharply bounded,” inasmuch as mathematicians “from time to time discover (or invent) new kinds of entities...which may be subsumed under the concept of number though distinct from any previously recognized sub-concept,” without treating “such additions to the extension of “number” as altering its meaning.” [Baker and Hacker, 1980, 351] But this much could also be said of, e.g., “protozoan,” given the evident possibility of discovering new species of protozoa; yet this word *does* seem to be one that biologists have given rigid limits (by means of an official definition), and hence ought not to be counted as a family resemblance term.⁷ Or, again: for centuries chemists discovered (and invented) new kinds of substances that may be subsumed under the concept of an organic compound, though distinct from any previously recognized organic compounds—simply because the concept of an organic compound had not yet been “recognized.” Yet, this does not seem to be the sort of concept that Wittgenstein would consider to be a family resemblance concept.⁸

Scarcely more helpful in making sense of Wittgenstein’s response to his interlocutor is Hallett’s gloss to the effect that “there is no fixed list of family characteristics, nor fixed number required for admission, nor sharp border for the individual characteristics themselves.” [Hallett, 1977, 150] Other commentators think of a family resemblance kind as essentially one for which the cri-

teria for membership are somehow linguistic, and assume that the adequacy of Wittgenstein's reply to his interlocutor follows straightaway.⁹ But the interlocutor's objection is not that there exists some *non-linguistic* condition (or, set of conditions) necessary and sufficient for gamehood—it is just that there *is* a necessary and sufficient condition for gamehood—linguistic or not (see [Grandy, 1979]). Anyhow, there is no non-linguistic set of conditions necessary and sufficient for falling under the concept of being a word of Esperanto, but that does not make this a family resemblance concept; this is so even if we could never be in a position to give a complete enumeration of everything that falls (much less has fallen or will fall) under it. This latter sort of consideration also casts doubt on Bambrough's construal of Wittgenstein's reply, according to which it amounts to the observation that one cannot reasonably aspire to “complete the enumeration of all” members of a family resemblance kind, since a family resemblance term “can be applied to an infinite number of instances without suffering any change of use.” [Bambrough, 1960, 219] But that much can be said of many expressions, such as “nephew,” “yesterday,” and “protozoan,” that seem not to fit the characterization that Wittgenstein has in mind for a word like “game” or “number.”

The first clue to a better development of the standard interpretation is contained in Wittgenstein's subtle shift from talk of kinds, properties, and concepts (e.g., of games and numbers) in §66 and §67 to talk of *words* (“game,” “number”) in §68. This suggests that Wittgenstein's thesis in these key paragraphs is a specifically *semantic* one, and I believe that only by carefully developing it as such shall we be able to appreciate the full force of Wittgenstein's reply to his interlocutor.¹⁰

As I understand it, family resemblance is a kind of literal content variability—a variability, that is, in what a term literally expresses or conveys from one use to the next.¹¹ What sets family resemblance terms apart from other kinds of variable content expressions is not that they exhibit some basic semantic property peculiar to themselves; rather, it is the fact that they combine a number of semantic properties that are not coinstantiated elsewhere.

Certainly, if family resemblance is a form of content variability, it must be a very unusual form—at least, unusual by existing standards. Wittgenstein obviously thinks that various philosophical terms—expressions such as “understands,” “knows,” “causes,” “conscious,” etc.—are family resemblance expressions, and it is equally obvious that he thinks that the apparent intractability of many philosophical problems is rooted in this fact. This means that family resemblance cannot, for example, simply be a kind of ambiguity, since ambiguity is a readily detectable form of content variability, and so not likely to cause any serious philosophical trouble.

A different form of content variability is indexicality; i.e., *monosemous* content variability.¹² We know that Wittgenstein was aware of the phenomenon of indexicality; this comes out in the following passage from his *Remarks on the Philosophy of Psychology*:

Is it a contradiction if I say: “This is beautiful and this is not beautiful” (pointing at different objects)? And ought one to say that it isn’t a contradiction, because the two words “This” mean different things? No; the two “This’s” have the *same* meaning. “Today” has the same meaning today as it had yesterday, “here” the same meaning here and there. It is not here as with the sentence “Mr White turned white.” [Wittgenstein, 1980, §37]

However, if various philosophical terms are indexicals (and I believe that they are, on Wittgenstein’s view), they must be quite different from commonly recognized indexicals such as “this” and “today,” since these too are readily identifiable as expressions of variable content.

One class of content variable expressions that are somewhat less easily recognized as such are the polysemes—terms like “light,” “high,” and “green,” which behave somewhat as ambiguous expressions, but, so to speak, with more cohesiveness, such that it seems that their use to express different contents is not merely a lexical accident (as appears to be the case with a term like “bank”). In fact I shall argue that it is partly by virtue of sharing a certain property with polysemes—what I call “topicality”—that family resemblance expressions are hard to identify as terms of variable content (see §2.3, below).

In order for a topical expression to qualify as a family resemblance term, it must also exhibit a property I call “semantic openness,” which many non-family resemblance polysemes and indexicals also exhibit, but which has nonetheless received little serious attention from semanticists. I devote §§3–4 below to semantic openness. But we begin now with some development of the notion of topicality, and a defense of the claim that certain items of natural language are topical indexicals.

2. Topicality

Roughly speaking, by a “topical” expression I mean one that expresses different contents in different contexts of use, in such a way that *which* content it expresses in any given use depends on what is under discussion, in the context of that use. In developing the notion of topicality, I begin with some low- to mid-grade polysemes (“light,” “high,” and “green”), work up to the high-grade polyseme, “to get,” and finally concentrate on certain terms that I shall argue are topical indexicals (or else exceptionally high-grade polysemes).¹³

2.1. “Light,” “High,” and “Green”

Although there seems to be no generally accepted definition of polysemy, terms like “light,” “high,” and “green” are by any account good representative examples of polysemous expressions. Intuitively, the range of contents that each of these words can express has more cohesion than that which pertains to a straightforwardly ambiguous term, such as “bank” or “duck.” At the same time, it doesn’t seem as if “light” has quite the same meaning in the sentence, “I spent

the morning doing some light reading” as it has in the sentence, “Her suitcase was quite light.” Nor does it seem quite right to say that “light” expresses the particular content that it does in each of these sentences (i.e., some content related to ease of endeavor in the former, and to mass in the latter) by virtue of some meaning that it has in both sentences. At any rate, the situation with “light” is not as it is with an indexical, such as “today,” which clearly does express the various things that it expresses by virtue of having the particular meaning that it has. This semantic fence-sitting behavior (the fence here dividing ambiguity from indexicality) is the intuitive hallmark of polysemy. It would be good to have a clear account of what the fence-sitting really amounts to; but for now, I want to leave the notion of polysemy at this intuitive, unrefined level, and focus on a different semantic feature that sets polysemes apart from most commonly recognized indexicals.

Unlike indexicals such as “here,” “I” and “today,” polysemes depend for their contents wholly or almost wholly upon what is under discussion in the contexts in which they are uttered. It is this dependence of their contents upon such topical contextual features (i.e., features of contextual topic of conversation) that makes polysemes like “light,” “high” and “green” topical expressions.¹⁴

Definition: A *topical expression* is a variable content expression that depends for its content in any context of use upon features of what it is used to speak of, in the context of that use (i.e., on the “topic of conversation,” relative to that context).

It is easy to see that the words “light,” “high” and “green” fit this definition. In these sentences:

- “There was a light breeze from the south.”
- “This is a light package.”
- “The fire won’t last with such light fuel.”

—“light” expresses a content having to do with intensity, weight, and density, respectively. And this clearly results from the fact that in the first sentence, “light” is used to speak of a breeze, whereas in the second it is applied to describe a package, and in the third, to some fuel for a fire. Similarly, “high” expresses a property relating to altitude in “Those clouds are very high,” as a result of being used to characterize clouds, whereas it expresses a sound-related property in “You can play some very high notes on a clarinet,” as a result of being used to describe (musical) notes.¹⁵

The adjective “green” expresses a property having to do with inexperience in the sentence, “These soldiers are still pretty green,” since it is used here to speak of human beings (in a professional vein). To find out what it expresses in the sentence, “These apples are green,” however, we have to look beyond the sentence itself to the broader conversational context:

M: Are we picking apples this weekend?

N: No, they're still green.

—versus:

M: What color will these apples be when they're ripe?

N: Those are Granny Smiths; they'll be green.

But here as in the other cases, the work of fixing a content for the polyseme is done by topical contextual features—features of what is being discussed, in a given context.

In contrast to polysemes, the indexicals “I” and “today” are what we may call “situational” terms, which depend for their contents on contextual features other than topic of conversation—features such as the physical location of a context (“here”), who is speaking in it (“I”), and when the use takes place (“today”).¹⁶ Naturally, such expressions help to *determine* what is under discussion whenever they are used, as indeed almost any referring expression does. But, for example, the content of “I” as uttered by me on a given occasion in no way depends on what I use it to say of myself: whether I am speaking of my health, my work, my ideas, or my relation to other people, places or things, “I” refers to me, in any literal use that I make of it. Similarly, in every literal use, “today” refers to a particular day, regardless of what is said about it (e.g., that it is hot, or dreary, or one that will live in infamy).¹⁷

2.2. “To Get”

The polysemes “light,” “high,” and “green” are rather low grade, bearing a distinct resemblance to ambiguous terms. In particular, like ambiguous expressions, these polysemes pretty much wear their content variability on their sleeves. Now I want to consider a polyseme that behaves more like a monosemous term, and that is somewhat less obviously an expression of variable content: the verb, “to get.”

We use this verb in a variety of general ways, some of them illustrated by the following sentences:

“I need to get some eggs.”

“It’s getting cold.”

“I haven’t got any money.”

In the first sentence, the verb “to get” is used in what we might call its acquisitional sense—a sense in which it means something similar to (but not, as will become evident, the same as) “to acquire.” In the second sentence, the verb is used with a meaning similar to that of “to become,” and in the last it simply occurs as a past participle of the verb “to have.”

I want to focus on the first, acquisitional use of “to get.” Consider it as it occurs in the following sentences (imagining appropriate conversational settings, as required):

“I need to get some eggs.”

“You are going to get a hangover.”

“He got the measles from his cousins.”

“I got a car at the police auction.”

“Did you get a confession?”

“I’m getting a bad feeling about this.”

“I got twenty years, fifteen suspended.”

“Try to get the waiter’s attention.”

“John got a severe reprimand.”

“We got three mallards, a black duck, and a canvasback.”

“They got his signature.”

“I got this idea from Aristotle.”

That “to get” varies in content across these uses is suggested by some substitutional considerations. In the sentence, “He got the measles from his cousins,” “got” may be substituted by “contracted,” *salva* truth-conditions. But such a substitution fails in the case of “I need to get some eggs” (“I need to contract some eggs”—?). Similarly, while “get” may be substituted by “suffer from” in the sentence, “You are going to get a hangover,” it may not be so substituted in: “Did you get a confession?” And various other cases of substitution failure can be generated, using the sentences in our list.

Such substitutional discrepancies alone do not show that “to get” varies in content. It may be that the verb simply expresses a complicated “disjunctive” relation, such that it could be defined something like this:

x gets $y \leftrightarrow x$ acquires y or x suffers from y or x contracts y or x purchases y or x elicits y or x begins to have y or x is condemned to y or x attracts y or x is subjected to y or x shoots y or ...

—at least, the verb’s expressing such a complex relation would be consistent with the foregoing substitutional observations. However, other considerations indicate that “to get” does not express this kind of disjunctive relation, but rather varies in content, now expressing one simple dyadic relation, now another. If the verb expressed the proposed disjunctive relation in every literal use, then any pair of things satisfying that relation would be such that it would be true to say of one that it got the other. But, while “John was subjected to a variety of annoyances” might express a truth, “John got a variety of annoyances” does not. This discrepancy could not arise if “got” invariably expressed the proposed disjunctive content. Likewise the disjunctive account cannot explain how “John came to suffer from the incessant clamor” might express a truth, even though

“John got the incessant clamor” does not. To take just one more example: “Smith shot a tin can” might express a truth simultaneously as “Smith got a tin can” expressed a falsehood. These examples and others that it is possible to bring forward show that “to get” does not invariantly express a complex disjunctive relation, but rather varies in content, sometimes expressing one dyadic relation, sometimes another.

Like the low-grade polysemes “light,” “high,” and “green,” “to get” is topical, depending for its content in a given context of use predominantly on what is under discussion, relative to that use.¹⁸ This dependence is ordinarily quite straightforward: in a typical use, the relevant topical features are determined by the expressions that fill the argument places in “*x* gets *y*” (or, “*x* got *y*,” etc.). For example, in the sentence, “John got the measles,” the relevant topical contextual features are John and the measles, or more exactly, the ordered pair, <John, the measles> —features that are determined by the presence of the expressions “John” and “the measles,” in that order. In the sentence, “John got the measles,” therefore, these expressions do double-duty: they help determine what is under discussion, and at the same time contextually determine the topical features that fix the content of “got” (which then in turn makes its own contribution to the determination of what is under discussion).

Although the topical features that determine literal contents for “to get” are normally ordered pairs determined by the syntax of the sentences in which the verb is used, there are cases in which the relevant topical features are not so determined. For one thing, there are cases in which one or both of the argument-places accompanying the verb is quantified into, such as in:

“Did John get anything?”

Here, syntactic features of the sentence in which “get” occurs are not enough by themselves to determine all the topical features needed to fix a content for the verb (recall the case of “green,” above, p. 489). Whether (e.g.) John’s having made some purchase, or alternatively, his having caught a fish, or elicited a confession, or contracted an illness (etc.) would warrant a positive answer to the question depends on the general topic of the conversation in which it was posed. Topical contextual features can also determine a content for “to get” in this relevance-mediated way in cases where the verb *is* flanked by referring expressions. Consider the following sentence:

“Smith has gotten the virus.”

In a normal conversational setting, this sentence would express a proposition to the effect that Smith contracted a certain virus. But one could imagine an unusual conversational setting in which this was not so. Perhaps Smith is part of a biological research team including Jones and Brown, and it is Jones who

utters the sentence to Brown, who has been expecting to learn the outcome of Smith's recent journey to the Congo to gather samples of a rare pathogen:

Brown: Have you heard from Africa?

Jones: Yes: I'm pleased to report that Smith has gotten the virus, and is shipping it to us as we speak.

In this conversation, "gotten" must express something like "acquired" or "successfully collected a sample of," in order for Jones's contribution to the conversation to be relevant. A more thorough discussion of these matters would scrutinize the norm of relevance itself, to see how it can be understood in terms of other conversational norms and commitments; here I have only tried to indicate some of the general ways in which relevance can help to fix a literal content for a topical expression.¹⁹

2.3. Topical Indexicality

Although the content variability of "to get" is less manifest than that of lower-grade polysemes such as "light" and "green," it is still not terribly difficult to detect (at least, once one knows what to look for). If philosophical terms—terms like "knows," "conscious," and "causes"—vary in content, they must do so in a yet stealthier manner for their content variability to have gone undetected for so long. We have seen that content variability becomes increasingly harder to detect in the case of a topical expression, the closer that expression approaches to monosemy. This suggests that if philosophical terms are topical expressions, they must be topical *indexicals*; or, at the very least, they must lie on the border between high-grade polysemy and indexicality.

Wittgenstein introduces the concept of family resemblance using as his paradigm the German word "*Spiel*," which translates into English as "game." Although I do believe that "game" can be shown to be a topical indexical, I shall be working with some alternative examples here: the English expressions, "to damage," "to be bad for," and "purpose."²⁰ The main value of using "game" to illustrate the notion of topical indexicality (and family resemblance) is that it does not *seem* on casual examination to exhibit these features, which (again) is important, since if topicality and family resemblance are to be of consequence to philosophy, that can only be because they are features of various *philosophical* terms that (like "game") do not wear it on their sleeves. Since the terms I have selected resemble "game" in this respect, they are appropriate substitute exemplars.

In order to show that "to damage" is a topical indexical, we must begin by showing that it varies in literal content. We can establish this using the same kind of argument that we used in the case of "to get." Consider the following uses of the verb:

- (1) "As you can see, the vandal damaged that painting."
- (2) "Be careful not to damage that shard: we want to study it later."

- (3) "I'm afraid her testimony has damaged our case considerably."
- (4) "This drought is really going to damage the crops."
- (5) "A vicious rumor damaged his reputation."
- (6) "Low grades have damaged my prospects for getting into law school."
- (7) "Alcoholism has already begun to damage your liver."

Now we note substitutional discrepancies suggesting that the verb is used in these sentences with a variety of contents. For example, in the sentence, "A vicious rumor damaged his reputation," "damaged" may be substituted by "diminished," *salva* truth conditions. But such a substitution fails in the case of "Be careful not to damage that shard" ("Be careful not to diminish that shard"—?). Similarly, while "damaged" may be substituted by "weakened" in the sentence, "I'm afraid her testimony has damaged our case considerably," it may not be so substituted in: "The vandal damaged that painting with spray paint." And various other cases of substitution failure can be generated, using the sentences in our list as well as others containing the verb "to damage."

As before, such substitutional discrepancies do not by themselves show that "to damage" varies in content. It may be that the verb simply expresses a disjunctive relation, such that it could be defined something like this:

x damages $y \leftrightarrow x$ diminishes the appearance of $y \vee x$ makes y less suitable for a certain intended purpose $\vee x$ weakens $y \vee x$ negatively affects the health of $y \vee \dots$

But, again as before, other considerations indicate that "to damage" does not express this kind of disjunctive relation, but rather varies in content, now expressing one (simple) dyadic relation, now another. If the verb expressed the proposed disjunctive relation in every literal use, then from "The drug will diminish your pain" we should be allowed to infer: "The drug will damage your pain"; but this inference is invalid. Nor can the disjunctive account explain how "They weakened the requirements for club membership" might express a truth, even while "They damaged the requirements for club membership" did not. To take a final example: one's appearance may well be diminished by a bad haircut or ill-fitting suit, without one's thereby being damaged.

The foregoing considerations show that "to damage" does not invariantly express a complex disjunctive relation, but rather varies in content, sometimes expressing one simple dyadic relation, sometimes another. It is no use objecting that "to damage" does not vary in content, on the grounds that it can be defined (e.g.) thus:

x damages $y \leftrightarrow x$ affects y in a way that is bad for y (without completely destroying y).

The problem is not that this analysis is incorrect: it seems to me that it *is* correct. But far from showing that "to damage" is content invariant, the correct-

ness of the analysis just shows that one or more of the terms of the *analysans* varies in content too. It is as if one were to argue for the content invariance of “today,” by pointing out that this word can always be substituted by the expression, “the present day,” *salva* truth conditions.

In fact, both “to affect” and “to be bad for” are topical indexicals; here I’ll focus on the latter phrase. Consider it as it occurs in the following sentences:

- (1) “Smoking is bad for you.”
- (2) “Smoking is bad for your heart.”
- (3) “Too much sunlight is bad for ferns.”
- (4) “Humidity is bad for period furniture.”
- (5) “Frequent rapid acceleration is bad for the engine.”
- (6) “High interest rates are bad for corporate earnings.”
- (7) “High unemployment is bad for the economy.”
- (8) “Pollution is bad for the environment.”

Sentence (1) may be paraphrased as follows: “Smoking reduces your ability to live as you would if you were healthy”; (3) admits of a similar paraphrase. Sentence (2), however, cannot be so paraphrased (“Smoking reduces your heart’s ability to live as it would if it were healthy”—?); rather, it must be paraphrased something like this: “Smoking reduces your heart’s ability to *function* as it would if it were healthy.” So, while it seems true that in general, if x reduces y ’s ability to live as y would if y were healthy, then x is bad for y , the converse at least does not hold.

In (4), “is bad for” may be replaced by “reduces the aesthetic, material, and/or practical value of” (although this substitution does not work for (1), (6), and (7)). But it is possible for something to reduce the aesthetic, material, and/or practical value of a thing without being bad for it. Shaving your head may reduce your aesthetic value without being bad for you. Washing Mickey Mantle’s autograph off of a baseball reduces its material value, but it isn’t bad for the ball. Losing the combination to a safe reduces its practical value, but it isn’t bad for the safe.

These last considerations already undermine a disjunctive approach to the analysis of *being bad for*; i.e., an analysis such as:

$$x \text{ damages } y \leftrightarrow x \text{ reduces the ability of } y \text{ to live as } y \text{ would if } y \text{ were healthy} \\ \vee x \text{ reduces } y\text{'s ability to function as } y \text{ would if } y \text{ were healthy (or, in good} \\ \text{order)} \vee x \text{ reduces the aesthetic, material, and/or practical value of } y \vee \dots$$

Additional considerations also block the move to disjunction. In (6), “are bad for” may be substituted by “cause...to decrease” (to yield: “High interest rates cause corporate earnings to decrease”). But it is false that in general, if something causes a thing to decrease, then it is bad for it. The setting of the Sun causes the temperature to decrease, but it is not bad for the temperature. We

may paraphrase (7) as follows: “High unemployment affects the economy in a way that diminishes the economic well being of the people who live in it.” Similarly, (8) may be paraphrased thus: “Pollution affects the environment in a way that diminishes the health of the things that live in it.” (However, we may neither replace (7)’s “economic well being” with (8)’s “health,” nor *vice versa*.) But it is false that in general, if something affects x in a way that diminishes the economic well being (or, health) of the people (or, things) that live in x , then it is bad for x . For example: fumigation diminishes the health of the things that live in my basement (cockroaches, mainly); but it is not bad for my basement.

In light of all these observations, we may conclude that “to be bad for,” like “to damage,” varies in content. Other apparently correct analyses of “to damage” suggest the content variability of yet other expressions. E.g., the following analysis seems correct:

x damages $y \leftrightarrow x$ diminishes y 's ability to serve one or more of y 's purposes.

—and this suggests that the term “purpose” figuring in the *analysans* is an expression of variable content. To corroborate this, consider the following sentences:

- (1) “As a doctor, my purpose is to make people healthy.”
- (2) “The purpose of one’s heart is to circulate blood.”
- (3) “This table is serving various purposes at the moment; among other things, I’m using it as my desk.”
- (4) “The purpose of the engine is to make the car go.”

In sentence (1), we may substitute “purpose” with “value to society” (perhaps with appropriate grammatical modifications to the rest of the sentence)—a substitution that is not possible in the other three sentences. (3) admits of the following paraphrase: “This table is being used in various ways at the moment; (etc.)”; but no similar substitution is possible in (1) or (2). (“The heart is used to circulate blood”—? Used by what?) And the move to disjunctive analysis is blocked by such facts as that a shoe can be used as a hammer without having hammering among its purposes, and that a person can be of value to society as a blood donor, without such donation’s being his purpose, in any capacity.

That the expressions “to damage,” “to be bad for,” and “purpose” are topical is clear enough, once it becomes clear that they vary in content; e.g., it is clear that the first verb expresses a different content in “Her testimony has damaged our case” from what it expresses in, “The drought has damaged the crops,” not because of some local feature of the contexts in which the word is uttered in each case, but because of its being used to speak of different things (the effects of some testimony on a law case versus the effects of a drought on some plants). The verb would express what it expresses in the first sentence even if it

were uttered in the middle of a desiccated cornfield, and it would express what it expresses in the second even if it were whispered by an attorney to a worried client. And we witness the same sort of responsiveness of content to topic in the case of “to be bad for” and “purpose.”

But why should the topicality of these expressions serve to disguise their content variability? I suggest the following explanation. Unlike situational indexicals such as “here” and “today,” predominantly topical expressions do not typically render the phrases in which they occur themselves content variable. This is true both of polysemes and predominantly topical indexicals. For example, the occurrence of “here” in the sentence:

“The cupcakes are here.”

renders the entire sentence indexical. By contrast, the sentence

“The cupcakes are green.”

is *not* rendered indexical by the occurrence in it of “green.”²¹ Predominantly topical expressions simply tend not to render the expressions (phrases or sentences) that contain them expressions of variable content.²² This suggests why it is comparatively difficult to identify predominantly topical expressions as terms of variable content: it is because we are seldom confronted with such a term without simultaneously being provided the contextual features required to fix a content for it. Thus when it comes to predominantly topical terms, it tends to be true that “what you see is what you get,” just as is the case with terms of invariable content. By contrast, this is not the case for a situational expression. What is the content of “here” in a present utterance by me of the sentence: “The cupcakes are here”? You don’t know, because you don’t know where I am: unlike “The cupcakes are green,” the sentence “The cupcakes are here” does not by itself give you the contextual feature (location of utterance) that fixes a content for the content variable term occurring in it.

Of course, the fact that an expression is predominantly topical does not by itself make it difficult to identify as a term of variable content. Low-grade polysemes are predominantly topical, but it’s not hard to see that they vary in content. But when a predominantly topical term approaches nearer to monosemy (such as “to get” does), it may be harder to identify it as a variable content expression. And when such a term *is* monosemous, its content variability may be very hard to detect indeed, since here we have neither situationality nor resemblance to ambiguity to alert us.

The fact that expressions like “to damage” and “purpose” are monosemous (or, nearly so) in addition to being predominantly topical makes them very difficult to identify as terms of variable content, and I believe it is this that accounts for our initial resistance to the suggestion that their contents do, in fact,

vary. But are these terms *monosemous* (and therefore indexical), or merely *nearly monosemous* (and so merely very high-grade polysemes)?

The question is hard to answer, mainly for lack of professional consensus on the notion of polysemy, which is a technical concept still largely up for grabs. If we follow Bréal (who introduced the term “polysemy”), we should almost certainly not count a term like “to damage” as polysemous. Bréal’s examples of polysemes are of such low grade as to be borderline ambiguous: “base” (as in “the base of the pedestal,” “military base,” and “acid and base”), “key” (as in “lock and key” and “singing off key”), and “act” (“a play in three acts,” “the Stamp Act”). [Bréal, 1897, 154–56] On the other hand, Eva Kittay uses “polysemy” to refer roughly to what I have in mind when speaking of topical content variability.²³ On this understanding of polysemy, a term like “damage” is polysemous, but may for all that be monosemous (in the sense in which I have been using the term “monosemous”)—or at any rate, such an understanding does not rule out a compatibility of “polysemy” with monosemy.²⁴ However, it may be worth noting that Kittay’s ingenious “Amelia Bedelia” test seems to ferret out—not topicality *simpliciter*, but—only what I (but not Kittay) call “polysemy.”²⁵ For, while we can construct Bedelia scenarios for terms such as “light,” “green,” and “to get” (as well as brutally ambiguous terms, such as “draw”)—we ask Amelia to get the butler, whereupon she shoots him—it does not seem possible to construct such scenarios for an expression like “to damage” or “purpose.” This suggests that “damage,” “purpose,” etc. are not (in my sense) polysemous, but lie rather at the monosemous end of the spectrum of content variability.

This result would comport well with Wittgenstein’s remarks on family resemblance, which certainly suggest that he would have rejected construing family resemblance as a form of ambiguity (even a “fancy” form, like polysemy). This in turn lends support to an interpretation of family resemblance partly in terms of topicality. From a purely philosophical standpoint, however, it probably doesn’t matter much whether we categorize terms like “damage” and “purpose” as monosemous or polysemous. The important point for philosophical purposes is that there are predominantly topical expressions that are either monosemous or close enough to monosemous that it is easy to mistake them for terms of invariable content. In the end, it is the possibility of such a mistake in the case of a term like “knows,” “conscious,” or “causes” that makes topicality a semantic property of potential interest to philosophy.²⁶

3. Essentially Open Practices

So far we have established the topical indexicality (or, near indexicality) of a variety of expressions. This alone is enough to overcome an objection (such as the one that Wittgenstein’s interlocutor raises in §68) to the effect that the concept of damaging or of a purpose can at least be defined in terms of some disjunctive condition necessary and sufficient for anything’s damaging anything

else (or, itself), or for anything's being (or, having) a purpose. But, even so, it might be possible to formulate a deepened version of the interlocutor's objection that takes the topical indexicality of such terms as "damage" and "purpose" in stride. For, even if it is impossible to "give the boundary" of the literal use of such a term by means of a rule that specifies the same content for it in every use, still, it might be possible to give such a boundary by means of a rule that, while not fixing the *same* content for the term in every use, nonetheless fixes *some* content for it relative to each use. That is: even though we have found it impossible to give a strict definition of (e.g.) "to damage" of the sort that it is possible to provide a term like "prokaryote," it may yet be possible to circumscribe the use of the verb by means of a rule of the sort that (strictly) determines literal contents for a term such as "today."

In this section and the next my purpose is to show that not even this much is possible. I shall argue that various terms, including the topical indexicals we have already examined, exhibit a characteristic I call "semantic openness," which makes it impossible to describe their literal use correctly through the precise specification of *any* strictly content-determining rule. In fact, many common practices are characterized by a certain kind of essential open endedness, and it will help if we approach the phenomenon of semantic openness first by identifying essential open endedness in some non-semantic cases; this is the task of the present section. In §4, I develop the notion of semantic openness in particular, argue that "damage" and other expressions are semantically open, and explain how this provides the resources to counter the interlocutor's objection, even in its deepened form.

3.1. Common Law

Consider our common system of law. To try to render this system as a body of strict rules would be to destroy its essentially "common" nature, and to replace it with a statutory system of law. It would also entail sacrificing many desirable features of our legal system, such as its ability to accommodate new and unforeseen cases in a prudent and reasonable manner. The practice of judicial or juridical decision in such a legal system is essentially open ended; for, although there exist general (e.g., constitutional) rules governing what decision a judge or jury may return in a given case, it may (and often does) happen that a case arises concerning which a decision cannot be made solely on the basis of existing statute, or statute combined with precedent. *Vis à vis* such a case, the judge's or jury's decision must, within constitutional limits, set a precedent without relying on any. Thus in common law the judge or jury often exercises a certain amount of *discretion* in reaching a verdict. The amount of discretion depends on the case, and in some cases may even vanish to a single option, such as when strict liability laws force a particular verdict. More often, however, statutes in a common law system leave judges and juries free to reach any of a number of legally admissible verdicts. This makes the practice of judicial decision in common law essentially open ended (in contrast to its "closed" statutory counterpart).

3.2. *Jazz*

The rules we have appealed to in characterizing the distinction between common and strictly statutory systems of law are simply legal regulations—laws encoded in a constitution or enacted by a legislative body. In characterizing the related distinction between improvisational and unimprovisational musical styles, we must appeal to rules of a different sort. These rules usually come in the form of a musical score—a set of instructions on how to play (which notes to produce, and when and how). A closely scored piece leaves the players no discretion as to how to play at any given point; playing such a piece is therefore what I call a closed practice.²⁷ By contrast, musicians playing improv do have discretion—sometimes an enormous amount of it—as to what musical act to perform at any given stage of the number. As with judicial decision in common law, the amount of discretion that the score allows a given jazz improviser may vary from one part of a piece to the next. For example, during his solo it might give him just a few bars to build on as he pleases, but require him to play in a very specific way at all other times. The more improvisational the piece, the more discretion the score will allow each musician at any given time. In extreme cases, there will be no score at all to provide guidance. But even in such a case, the musicians' playing will usually work within some kind of structure. However, the rules that structure the improvisation now are not a score, but something more mysterious: unspoken rules of phrasing and harmony that the players would probably be hard pressed to articulate (but quick to detect violations of).

3.3. *Conversation*

Another practice that exhibits essential open endedness is ordinary conversation. Although seldom made explicit, there are rules governing a conversation, placing constraints on how it may proceed at any given juncture; if these rules are violated, conversation ceases, perhaps degenerating into an exchange of verbal abuse or mere babbling. For instance, someone who makes an assertion during the course of a conversation is required by conversational rules to respond to legitimate challenges, and, if he cannot meet such a challenge, the rules of conversation oblige him to retract his assertion. By the same token, as long as his assertion fails to raise any unmet challenge, it is by conversational rule entered into that conversation's "common ground."²⁸ Thus, given the rules governing conversational practice, an assertion "creates a space of legal illocutions" for all parties to a given conversation (similar spaces are created by other conversational moves, such as the raising of a question).²⁹ But for all the rules of conversation, much of what we do in conversing is not mandated by conversational rule. The rules of conversation provide guidelines—sometimes broad, sometimes narrow—as to how we must proceed as a conversation develops, but ordinarily there are indefinitely many ways in which we can follow these guidelines. As Mitch Green observes in this connection,

we...need not assume conformity to conversational requirements must leave a speaker with at most one option. Where a conversation permits a choice we need say only that the conversation requires that at least one of the admissible options be taken... So long as the set of admissible options does not include all possible speech acts the injunction to do what the conversation requires will still be substantive. Which of the admissible options is appropriate to take may then depend...on factors, such as the speaker's epistemic state, that are external to the conversational requirements. [Green, 1997]

Indeed, it would be self-defeating to try to supplement the actual rules of conversation with further rules serving not as guidelines but as detailed specifications of what to say or do at each stage. The most we could accomplish by attempting this would be to codify rules of a linguistic practice distinct from conversation, and perhaps similar to litany.

As in the examples of common law and jazz, the rules governing conversation allow conversationalists a certain amount of discretion in what to do—i.e., in what speech-act to perform—at the stages at which they are called upon to act (in the course of conversing). Also as in the previous examples, the governing rules may leave an individual fewer options in some cases than in others, and in the limiting case, may even require a specific speech-act.³⁰ But ordinarily, the rules of conversation give conversationalists considerable powers of discretion over how to satisfy their dialectical obligations, which is what distinguishes conversation from even the most elaborate litany. The main difference between the legal and conversational cases is that the rules whose relative lack of stricture makes conversation an open practice are not statutes, nor for the most part explicit rules of *any* kind (such as a written score). They are rather rules of the sort we call “social norms,” and their force does not derive from any act of legislation. It is true that the rules governing a litany may be religious statutes, but this is not necessary; nor of course is it necessary that the rules of conversation *not* be explicit or even that they fail to carry the force of law. But in the legal examples, at least, the rules must have legal force, and this shows that the open/closed distinction in those examples is based on a distinction between different kinds of rules from that upon which the same distinction is based in the cases of linguistic exchange.

3.4. Analysis of Essential Openness

So far we have identified essential openness by contrasting various pairs of practices. Essential openness is what distinguishes common law from statutory law, improvisational jazz from closely scored musical styles, and conversation from litany. But to identify essential openness is not yet to gain clear insight into its nature, and this is what we must try to do now.

One thing that emerges from consideration of the foregoing examples is that the distinction between essentially open practices and their closed counterparts does not correspond to a straightforward distinction between *kinds of rules*—

e.g., between informal social norms and legal regulations, or between explicit and implicit rules. In some cases one kind of rule is especially appropriate or even essential (as in the legal cases, where the rules must be laws), but no such kind of rule is essential to *all* open (or, closed) practices, and in the case of many practices—open as well as closed—it is inessential what particular form the rules take (i.e., whether of norms, laws, written instructions, etc.).

Nor is the difference between closed and open practices a difference pertaining to the existence or nonexistence of some kind of (mathematical) *function* in the case of practices of each kind. For *any* given practice, open or closed, there is a function that assigns some action to each stage of the practice, which is the action that those who properly engage in the practice perform (or did, will, or would perform) at that stage. In the case of law, this is a function from cases to verdicts; in the case of practices of linguistic exchange, it is a function from stages of verbal exchanges to speech acts; and in the musical cases, it is a function from parts of pieces to notes.

The difference between open practices and closed ones is simply that the former but not the latter allow their practitioners some amount of discretion in engaging in them; for this reason, the sorts of functions referred to in the preceding paragraph underdetermine the practices to which they pertain. To see why, consider again the example of our common system of law. Suppose that a newly founded nation were looking for a legal system on which to model its own, and after doing some research decided that it wanted to adopt our (U.S.) legal system wholesale. It could not do so simply by copying existing American statute word for word. For if this new state were to copy our statutes and appoint judges whose job it was simply to determine whether the statutes applied, and (when they did) make the prescribed rulings, it would be adopting a strictly statutory legal system distinct from our common system of law. Indeed, even if the nation were to adopt a strictly statutory system of law that proved to be extensionally equivalent to our common system, in the sense that it just so happened that rulings in the two systems always coincided, its legal system would *still* be distinct from ours, since in our system but not in theirs judges and juries would typically exercise a certain amount of freedom in coming to their decisions. Yet if we were to describe each of these systems in terms of a function from cases to verdicts (however complex), the descriptions would match perfectly. Hence such a description would underdetermine either system, and would fail to capture something essential about each of them.

4. Semantic Openness

Before moving on to consider semantic forms of essential openness, it will help to consider briefly what the point is of having open ended practices. Why for example do we engage in the practice of conversation, instead of getting by with a fixed stock of ritual verbal exchanges? There are many reasons. To be sure, one of these is that a fixed stock of formulaic exchanges would have its range

of application limited in a way that the practice of conversation is not. But this is at most part of it; after all, if the fixed stock were quite large, its potential utility might extensively overlap with that of ordinary conversation. A deeper reason why we have conversation instead of some other, inflexible form of verbal exchange is out of a certain division of labor. Neither as individuals nor as a community do we see ourselves fit to anticipate the linguistic needs and desires of even all of our contemporaries, to say nothing of future language users. We simply—and, it seems, rightly—put more trust in the ability of the actual parties to a given conversation to see what moves will best further its aims than we put in our own ability to foresee, perhaps far in advance, what aims will arise, and what linguistic avenues would lead to their satisfaction. (Parallel considerations apply when weighing the advantages of common versus strictly statutory legal regimes.)

In addition, the essential open endedness of conversation is probably felt by many to be a basic good, as part of what makes conversation a thing to be valued for its own sake. This kind of consideration is perhaps more in evidence *vis à vis* the question as to why we have improvisational art forms. Yet here too we find a kind of division of labor at work, inasmuch as in the case of improvisational jazz, the work of composition is distributed among the playing musicians. And here again what motivates the division is a belief (sometimes mistaken) that the musicians will find musically interesting ways to develop a theme—ways that would not occur to any individual or team of individuals composing on paper, as opposed to “on the fly.”

Given the potential payoffs of essential openness, it would be surprising to find a practice as vital as our literal use of natural language entirely bereft of it. In developing the notion of openness in relation to semantic cases, we must first identify the relevant linguistic practice. This is the practice of using an expression to express contents, in a literal fashion (on various occasions). Using a term literally will count as an open practice just in case speakers exercise some amount of discretion in engaging in it; i.e., some discretion over what literally to express by means of that term, on at least some occasions.

It is clear that there are semantically *closed* expressions (i.e., ones the uses of which leave speakers with no discretion as to what literally to express by their means). For example, our use of the indexical “today” is closed, since it is governed by a rule that leaves us with no discretion over what to express by it, in any use (*viz.*, the rule that in every literal use, “today” expresses the function that correlates with each evaluation point the day on which that use occurs).³¹ If there are semantically *open* expressions (i.e., expressions the literal uses of which are essentially open practices), they will match up well with many of the things that Wittgenstein says about family resemblance terms. For example, the use of a semantically open term would not be “everywhere circumscribed by rules”—i.e., by rules dictating in a discretion-eliminating way what contents speakers may use them literally to express on any occasion. It would also seem appropriate to describe a semantically open term as one that we use

in such a way that its extension “is *not* closed by a frontier,” inasmuch as the exercise of discretion implies the possibility of setting precedent (more on this presently).³² The idea that family resemblance terms are semantically open also yields a useful interpretation of the Bambroughian schema:

ABCD ABCE ABDE ACDE BCDE ...

—for if semantic openness is a mark of family resemblance, we may take the ellipsis to represent the fact that the literal use of a family resemblance expression is open ended, in the manner characteristic of any essentially open practice.

To see that there really are semantically open terms, consider once more the verb, “to damage.” In claiming that this expression is semantically open, we are saying that its literal use leaves speakers with at least some discretion over what literally to express by it, in at least some contexts of use. Now, it is plain that our use of “damage” to express literal contents is not purely *ad hoc*, so we should expect to find some system of constraints governing the term in the same general, partial way as the rules of common law, jazz, and conversation govern those practices.

We have seen that those who engage in such practices can find themselves with varying degrees of discretion in the course of engaging in them, and in limiting cases with no discretion at all. But there is a further complexity here that we have not yet addressed. In the case of many open practices, the structuring rules are not laid down once and for all. They are formulated over time; in particular, they come into being through the establishment of precedent. The case of common law is an especially clear example of how this works. The constraints on judicial decision in common law are not imposed only by pre-existing statutes, but by precedent; i.e., by prior acts of judicial decision. How exactly this works is a matter of controversy, but the basic idea is that in rendering his verdict, a judge not only must rule in a way that is consistent with standing statute,³³ but also must make his ruling consistent with prior rulings (both his own and other judges’) on similar cases. Taking the role of precedent into account explains why it is often difficult to characterize the elements structuring an open practice. If the structure of such a practice is at least partly imposed by precedent, then it is at least partly a product of (antecedent) exercisings of discretion, as for example the constraints on judicial decision in common law in any given case are partly imposed by antecedent exercisings of judicial discretion.

So, with many open practices, we get a compounding effect in the constraining or structuring rules or elements: how people may exercise their discretion (as well as how much discretion they have) while engaging in such a practice at a given time depends at least partly on how others have exercised discretion in the same practice before them. When it comes to a practice as involved as the use of a common item of ordinary language, the body of constraining precedent in effect at a given time can be enormous, and the difficulty of describing the structuring constraints that govern such an expression may therefore be considerable.

Here, I shall simply try to indicate some of the ways in which precedent can structure the use of a semantically open expression, and does structure our use of the verb “to damage.” For this purpose, it will help to begin with a fictional case. Consider a language game involving players A, B, C, D, and E, who constitute a community that convenes on a monthly basis to propose and vote on decrees by which the members subsequently consider themselves bound. The language game (call it L1) begins when the five players meet in June, and (after some discussion) issue the following decree:

We hereby unanimously decree that

1. the word “damage*” shall in every use express a dyadic relation;
2. “damage*” shall express the relation of *diminishing the viability of* whenever an expression referring to a plant occupies its second argument place (i.e., the y-place in: “x damages* y”);
3. “damage*” shall express the relation of *reducing the utility of* whenever an expression referring to a tool occupies its second argument place;
4. what “damage*” expresses whenever anything besides a plant or a tool occupies its second argument place shall be established by a vote of at least a four-fifths majority of this committee.

Subsequent to this decree, the players put “damage*” to work in various ways; for example, B one day complains to D that locusts are damaging* his corn. Later, during their July meeting, C moves to add the following clause to the original decree:

- 3.1. “damage*” shall express the relation of *weakening* whenever its second argument place is occupied by an expression that refers to a case at law.

The motion is seconded, put to vote, and passes unanimously, and the utility of “damage*” is thereby increased. In the same way we can imagine further additions being made to the “damage*”-decree over time, resulting in further increases in the utility of “damage*.” We can also imagine the rejection of certain proposed additions—e.g., of a proposal to let the verb express the relation of *ridiculing* whenever one of its argument places contains a term that refers to a person. The players in L1 thus exercise their discretion over what to express by “damage*” by casting votes.

Moving closer to actual, natural language, consider a language game L2, much like L1, except that when the players convene in June, the decree they issue contains only three clauses—clauses (1.) through (3.) of the “damage*”-decree figuring in L1. Here we can again imagine C at the July meeting proposing an extension of the verb’s use to certain legal applications, the committee seconding and ratifying his motion, etc. That is, we can easily imagine a continuing series of official additions to the original (three-part) “damage*”-decree, along the same lines as we could imagine such additions being made in the case of

L1. In L2, we may say, the power of discretion sanctioned explicitly by the fourth clause of the decree issued in L1 is implicit in the players' adoption of further rules governing "damage*." However, as in L1, this discretion is exercised explicitly, by vote.

Now consider a game, L3, differing from L2 only in that in L3 the players cease to convene after their initial meeting (at which they issue the same decree as in L2). We can nonetheless imagine them using "damage*" in a continually expanding range of applications, just as they would were they mutually and explicitly to agree to each new kind of use (e.g., by a vote, as in L1 and L2). In L3, the players' exercise of discretion over what to express by "damage*" is itself embedded in their day to day use of the term.

To get from L3 to the actual use of our word "damage," we have only to discard the fiction of an original convention at which initial conditions for the use of the term get decided. We do not know what content(s) speakers initially used the verb "to damage" to express, but in all likelihood these initial uses were not the result of explicit agreement. However, *how* the verb began life is not as important as the manner in which it has matured, viz., by a process of precedent-setting improvisations, each constraining subsequent uses and at the same time providing additional structure around which for subsequent exercise of discretion to occur.

There are various ways in which a speaker might exercise discretion in the use of "damage" so as to set a precedent for future uses of it. For example, he might use the term metaphorically and, if the metaphor is apt, this way of using the term could eventually become part of its *literal* use. (This is not to suggest that semantic openness is somehow tantamount to dead metaphor. Even if the first application of "to damage" to describe the effects of legal testimony was a metaphorical use of the word, certainly *we* do not use "damage" metaphorically when we say of someone's testimony that it damaged the case for the defense.³⁴) Another way in which one might set a precedent in the use of "damage" would be by using it to express a new content while making it explicit that this is what one is doing, perhaps by saying something like this: "I'm afraid her testimony has damaged our case considerably, and by "damaged" here I mean "weakened"."³⁵ Or perhaps someone might simply utter the sentence "Her testimony has damaged our case," and set a precedent in the use of the verb by later explaining what he meant by it (perhaps to honor a request for clarification).

In many cases, a speaker exercises his discretion over what literally to express in using a certain term by having a certain intention, in using it. This happens most obviously in the case of certain demonstratives and ordinary indexicals, such as "that" and "here." The word "that" is governed by rules that at least loosely constrain what it can literally express, in a given context, and usually overt contextual factors (both topical and situational) take up the remaining slack to fix a specific content for the demonstrative, in a given use. But there are cases in which a speaker's use of "that" is such that there is a plurality of things that the speaker could reasonably be taken to be referring to

by it, even taking into consideration the rules that govern the word and the overt circumstances surrounding this particular use of it. To take a simple example: if I nod in the direction of two models who are standing next to one another, and say to you: “That woman is stunning!”, which content my use of “that” literally expresses may well depend ultimately on which of the two women I intended to refer to (with the demonstrative). An analogous case involving the indexical “here” might be the following: two graduate students are introduced to one another for the first time at a welcoming party for new additions to the philosophy department, and one of them asks the other: “How long have you been here?” Under the right circumstances, there may be no way to tell on the basis of overt contextual factors whether by “here” the speaker means “at this party” or instead, “with this department.” As in the scenario involving “that,” the speaker’s intention is the final court of appeal in determining what “here” literally expresses, in this context. This is not to say that the speaker’s intention overrides every other factor; it does not: the questioning student’s use of “here” would not have referred to the Solar system, even if his intention had been to use “here” to refer to the Solar system. The point is just that a speaker’s intentions are in some cases at least *partly* determinative of the literal content of a word like “here” or “that.”³⁶

The cases discussed so far provide clear examples of how speaker discretion can help determine a literal content for an expression (in a given context of use), thus establishing a precedent of some duration for the future use of the term in question.³⁷ Often, however, precedent in the use of an open expression is established through the agency of a gradual semantic drift, whereby a term repeatedly gets used to express new contents that are closely enough related to contents that the term has been used to express before that their novelty goes unnoticed. (It seems that it is this kind of phenomenon that Wittgenstein is getting at with his metaphors of “overlapping” and “criss-crossing”—see, e.g., §67.³⁸) It is not hard to imagine how this sort of thing might have happened in the use of “damage.” At one time it may have been common to use the verb to describe the effects of various natural forces on plants, such as in the sentence, “This mildew is damaging my wheat,” in which “damaging” expresses something like *diminishing the health of*. It would be a small and probably unnoticed step to go on to use the verb in such sentences as: “The mildew is damaging my fields,” wherein it expresses some such relation as that of *diminishing the health of what is growing in*. And another subtle shift might take us to: “The straw in the barn has been damaged by mildew.” In these examples, the successive new applications of “to damage” (to express new contents) overlap in that they relate to various stages of a single process (that of raising, harvesting, and storing crops).

In other cases, the overlap may result from similarities between the things that are said to do the damaging. In this way “to damage” might progressively be applied in the following ways, without anyone’s noticing that in each application, it is used to express a different content:

“Lightning damaged that tree.” [*affected so as to diminish the health of*]
 “Lightning damaged this monument.” [*affected so as to detract from the appearance of*]
 “Lightning damaged our antenna.” [*affected so as to impair the operation-ality of*]

Similarly, speakers accustomed to using “damage” to express a “utility diminishment” content, as in “We found a rust-damaged knife,” might easily go on to use it to express a content having to do with purely aesthetic diminishment, as in: “We found a rust-damaged medallion.” Here again, it is a similarity pertaining to what is said to do the damage (in this case, rust) that creates a natural bridge from one sort of application to the next.

When the use of “to damage” to express a new content is incremental in nature (such as in the foregoing examples), it sets a precedent with such a lack of fanfare that we are unlikely to realize that a precedent has been set (or, discretion exercised) at all. But it has been. Why does the verb express what it actually does when applied to fields, instead of some content having to do with, say, soil chemistry? Why do we say (truly) that aggressive tobacco farming damages not fields, but the soil? The natural explanation is that the precedents backing up the application of “damage” to fields had to do less with the ground than with what was growing in it. But any number of different kinds of application of the term might have been taken as precedents in extending its use to fields; in taking one particular kind of application as precedent for the application to fields, speakers exercised their discretion over how to use the verb.

To take another example: why should breaking a piece off of a wedding cake suffice to damage it only if the breakage occurs prior to the wedding? The aesthetic diminishment is the same, either way. The precedent for the use of “damage” in this case comes rather from applications in which it expresses contents related to the undermining of the intended use(s) of that which gets damaged (recall the case of the shard). And it seems natural to suppose that it is also due to the vagaries of precedent that one can poison a wedding cake (say, by carefully dusting it with an invisible cyanide powder) without damaging it—despite the fact that poisoning would undermine *one* of the cake’s intended uses. In extending the use of “damage” to speak of wedding cakes, speakers might have proceeded on the basis of any of a variety of antecedent uses of the term (or combinations thereof); in extending its use as they did, they exercised their discretion over what literally to express by “to damage” (in this case).

It would seem to be the precedent-setting exercise of discretion that explains the fact that it is possible to damage a painting (say, by vandalizing it) without reducing its aesthetic or material value (e.g., if it had neither kind of value to begin with). That it is possible to wash Mickey Mantle’s autograph off of a baseball without thereby damaging the ball (despite devaluing it considerably) is also to be chalked up to the effects of discretion in establishing precedent. And why, after all, isn’t losing the combination to a safe bad for it

(despite reducing its practical value considerably)? Here too we find evidence of discretion—presumably the precedent for applications of the expression “bad for” to safes came from uses of it to describe the effects of utility-diminishing physical alteration on other kinds of objects; as before, it was up to speakers to proceed on the basis of this particular precedent. Finally, in the case of “purpose,” we no doubt have the vagaries of precedent-setting discretion to thank for the fact that a shoe does not have hammering among its purposes, no matter how effective it might prove to be, thus employed.³⁹

Semantic openness comes in degrees, corresponding to degrees of speaker discretion in the literal use of a term, and only a closed expression is such that speakers never have any discretion over what literally to express by its means. Could an open term be synonymous with a closed one? Perhaps, if the open term were open only to an insignificant extent—i.e., if its use never left speakers with more than a modicum of discretion, and with none at all in nearly every application. But the openness of the terms we are considering is not insignificant, and no such expression can be synonymous with any closed expression: a difference as great as this between the literal uses of two terms implies a difference of meaning.⁴⁰

We can now further clarify Wittgenstein’s response to the objection his interlocutor raises in §68. Consider a semantic parallel to the legal thought-experiment we performed at the end of §3, above. Some people, let us imagine, are engaged in the project of designing a new language, from scratch (perhaps something like Esperanto). Among other things, they might want to include in their language a term with the same literal use as the English word, “today.” To that end, they could introduce a word (say, “today*”) by means of the following rule: relative to any use, “today*” is to refer to the day on which that use occurs. But they could *not* by means of any such rule introduce an expression with the same literal use as our word, “damage.” For any attempt to do so would yield a term regarding which speakers would have no discretion over what literally to express by its means, in any context; hence, even if the term—say, “damage*”—happened to coincide with our “damage” in terms of what it literally expressed relative to any context of use, it would still not have the same literal use—nor, therefore, the same meaning—as “damage,” since the use of “damage” but not “damage*” leaves a significant amount of content-determining work to be done by the exercise of speaker discretion.⁴¹

Thus, just as it would be a mistake to identify the essence of a legal system with a function from cases to verdicts (see above, p. 501), so too must it be a mistake to identify the meaning of an expression with its (Kaplanian) “character.” For as we have just seen, two expressions can share the same character while differing in meaning, as a result of being used as parts of very different linguistic practices. Identity of character is at most a necessary condition of identity of meaning.

This, then, is what Wittgenstein is getting at when he says that we use the word “game” in a way that is “not closed by a frontier”: his point is that this word has a meaning that cannot be specified simply through the specification

of a (mathematical) function (however complex), nor through the provision of *any* rule that would deprive speakers of all discretion over what literally to express by its means.

5. Conclusion

Given that it is tantamount to open topical indexicality, why is family resemblance a concept that should be of general interest to philosophy? The potential philosophical importance of family resemblance emerges when we consider what would be the consequences of the fact, if it were a fact, that, say, the verb “to cause,” or “to know” were, like “to damage,” an open topical indexical.

In general, it seems that the topicality of family resemblance terms stands to be of greater consequence to philosophy than their semantic openness.⁴² This is so especially if, like “to get” and “to damage,” a philosophically crucial term like “to cause” or “to know” has no “central” or preeminent content, among all of those that it has the potential to express.⁴³ The absence of any such content would preclude the possibility of adopting a “star it and get to work” response to the topicality of a given expression. For if among the contents that “to know” has the potential to express, none is in any important sense “central,” then no analysis of any one of its potential contents could qualify as a (good) analysis of *knowledge* (as opposed to some dull distant cousin, “knowledge*”).

A further way in which family resemblance stands to affect philosophy has to do with the method of counterexamples—i.e., of criticizing a general (philosophical) hypothesis by adducing a false instance of it. If very many philosophical terms exhibit family resemblance, then this critical method may not be as widely available as we now take it to be. A counterexample works only when it uses a term with the same content (and not merely the same meaning) that it has in the claim that it is meant to counterexample. Thus if a term varies in content without our realizing that it does so, we are liable to misapply the method of counterexamples (as well as its constructive correlate, the process of hypothesis refinement). For example: if “to know” is a topical indexical, then the fact that a false proposition is expressed when the verb is applied to describe a Gettier case may not yield a counterexample to the claim that to say of most normally situated people that they know that they have parents is simply to say that they have the relevant justified true beliefs. For if “knows” is topical, it may simply express a *different* literal content as used to describe a Gettier example from what it expresses when used to describe normally situated people’s epistemic relationships to their own parents.⁴⁴

All of this is at the level of conjecture. Whether some of the terms crucial to the formulation of philosophical questions and hypotheses really are topical indexicals (or, family resemblance terms)—and if they are, whether they lack “central” contents of the sort presupposed by the “star it” approach—can only be discovered through a close study of the semantics of philosophical language. That, however, is a task for another day.⁴⁵

Notes

¹See [Wittgenstein, 1958, §§66–68]. Unless otherwise indicated, section references are to [Wittgenstein, 1958].

²Witness the fact that family resemblance gets mentioned a total of six times in the new 500 page *Cambridge Companion to Wittgenstein*—each time in passing.

³For the *locus classicus* of the standard interpretation, see [Bambrough, 1960]. Other well-known representatives of the standard interpretation include [Fogelin, 1976, 117–123]; [Hallett, 1977, 140–50]; (in a critical vein) [Grandy, 1979]; and [Baker and Hacker, 1980, 325–32].

⁴See [Bambrough, 1960, 209–10]. In Bambrough’s notation, adjacent capital letters collectively stand for a conjunctive property or condition; so, for instance, something instantiates *ABCD* just in case it instantiates *A* and instantiates *B* and instantiates *C* and instantiates *D*.

⁵Grandy presses this objection in [Grandy, 1979].

⁶On this point I find myself in agreement with a number of commentators, including Fogelin (see [Fogelin, 1976, 118]).

⁷Compare [Wittgenstein, 1958, §18].

⁸That there are—or at least, were supposed by Wittgenstein to be—non-family-resemblance expressions is indicated by his remark that one *can* use a word for a “rigidly limited” concept (§68).

⁹For instance, J.F.M. Hunter suggests that “the criterion of class membership” for a family resemblance kind, *X*, is “linguistic: anything is an *X* if the linguistic community routinely so describes it.” [Hunter, 1985, 59]

¹⁰I say “carefully”: it will not suffice merely to adopt a free-and-easy style of exposition that uses terms like “word” and “concept” interchangeably. Here I give the following example from Fogelin, only because it is purely typical of the literature: “...do many of the concepts of our ordinary language function as Wittgenstein says they do?” [Fogelin, 1976, 121]

¹¹Here and throughout, I work with a notion of “content” derived from recent work in the philosophy of language. Roughly and intuitively, the content of an expression relative to a given use is what that expression is used to say, on the occasion of that use. More specifically, “content” is generally regarded as a function from evaluation points (such as possible worlds, world-time pairs, or more fine-grained “situations”) to individuals, classes or properties (see, e.g., [Braun, 1994, 193]). More precisely, we can define what it is for an expression token to express a given content, as follows.

Let “*e*” variables range over tokens of singular terms (such as names and pronouns), “*ϕ*” variables range over tokens of monadic predicates and common nouns, “*ρ*^{*n*}” variables range over *n*-adic predicate tokens (*n* > 1), and “*σ*” variables range over sentence tokens. Let “*I*” variables range over world-to-individual functions, “*F*” variables range over properties, “*R*” variables range over *n*-adic relations, and “*P*” variables range over propositions (functions from worlds to truth values). Let “*x*” variables range over individuals, and “*w*” variables range over possible worlds. Then we may define the relation of *expressing a content* by cases, as follows:

1. $(\forall \epsilon, I)(\epsilon \text{ expresses } I \text{ iff } (\forall w, x)(\epsilon \text{ refers to } x \text{ in } w \text{ iff } I(w) = x))$.
2. $(\forall \phi, F)[\phi \text{ expresses } F \text{ iff } (\forall w, x)(x \text{ is in } w \supset (\phi \text{ applies truly to } x \text{ iff in } w, F(x)))]$.
3. $(\forall \rho^n, R^n)[\rho^n \text{ expresses } R^n \text{ iff } (\forall w, x_1, \dots, x_n)(x_1, \dots, x_n \text{ are in } w \supset (\rho^n \text{ applies truly to } x_1, \dots, x_n, \text{ iff in } w, R^n(x_1, \dots, x_n)))]$.
4. $(\forall \sigma, P)(\sigma \text{ expresses } P \text{ iff } (\forall w)(\sigma \text{ is true in } w \text{ iff } P(w) = \text{Truth}))$.

(In speaking of what a token τ refers or applies to “in *w*,” I speak of what τ —situated in whatever world it happens to inhabit (which may or may not be distinct from *w*)—refers or applies to in *w*, not of what τ (or, a token of τ ’s type) would refer or apply to, if produced in *w*.)

In the text I also refer to something called “indexicality,” an indexical being an expression type that has the capacity to vary in literal content across a range of tokenings with respect to which it retains a single meaning (common examples of indexicals are pronouns such as “I” and “you,” and expressions like “here,” “today,” and “this”). The *meaning* of an indexical is often identified with a function that maps contents onto contexts of use; this is what David Kaplan calls a “character” (see [Kaplan, 1988]). However, in §4 I argue that meaning cannot be identified with character.

¹²I use the term “meaning” and its cognates with a broadly Strawsonian sense—a sense in which “[t]he meaning of an expression cannot be identified with the object it is used, on a particular occasion, to refer to,” nor the meaning of a sentence identified with “the assertion it is used, on a particular occasion, to make,” but rather with “the rules, habits, conventions governing its correct use, on all occasions, to refer or to assert.” [Strawson, 1990, 321] Although linguistic meaning under this conception bears a closer resemblance to Kaplan’s “character” than to his “content,” I do not identify meaning with character (again, see §4, below).

¹³In my terminology, the more similar a polyseme is to a straightforwardly ambiguous term, the lower its grade, and the more it resembles a monosemous term, the higher its grade.

¹⁴Generally speaking, I take “topic” to be largely a function of the conversational commitments that are in effect at a given stage of a conversation (e.g., as determined by Gricean norms). But for present purposes, we can make do with the notion of topic in a rough-and-ready unanalyzed state.

¹⁵Cf. §377.

¹⁶Recanati nods in the direction of the distinction between topical and situational content variability, but does not go on to develop it in any detail (see [Recanati, 1997, 3–5]).

¹⁷Topicality and situationality are not mutually exclusive, and various indexicals exhibit both properties, in varying degrees of combination. Take the case of “here.” The content of this indexical frequently depends on topical contextual features, as in the following conversation:

M: So, how do you like Alabama?

N: Well, you know, it sure is hot here.

—which may be compared with this conversation, in which topical features fix a different content for “here”:

M: How long have you been living in this building?

N: Oh, for quite some time. We like it here.

In other cases, “here” depends for its content more on situational contextual features. If I hand you a dollar with the words: “Here is a dollar,” the content of “here”—i.e., roughly: “(in) my hand”—gets determined independently of what, if anything, we were discussing at the time.

¹⁸While “to get” relies for its literal content predominantly on topical contextual features, it does not rely wholly on such features, in every context. In certain contexts its content may be partly determined by situational factors. For example: if we bump into each other as I’m leaving a pet shop carrying a cage containing what are obviously ducks, and I say to you: “Look, I got three mallards, a black duck, and a canvasback,” my use of “got” does not express what it would were I to utter the same sentence to you while carrying a shotgun and a brace of clearly dead ducks. Thus, in the pet shop scenario, an utterance on your part of “You got some ducks without shooting them” would express a truth, and not a contradiction. (Of course, the same thing could happen as a result of variant topical factors: e.g., “got” might not express the shooting content if uttered in the course of a discussion between us of my recent trip to the pet shop. This shows that the sketch given in the text of how “to get” varies in response to topical contextual factors is rough and incomplete; unfortunately, filling it in would take up too much space, here.)

¹⁹The effect of the maxim of relevance in the cases we have been considering is that of fixing (or, helping to fix) a literal content for “to get,” rather than the pragmatic one of generating an implicatum. There is no need to assume that whenever a Gricean conversational maxim comes into play, its effect must be that of generating such an implicatum. In general, a maxim can play that role only given that the expression in question (e.g., the verb “to get”) already expresses a certain literal content. But what is this content, and what are the Gricean maxims, that combine to generate the content conveyed by “to get” in each of our examples? What, short of the relation of contracting, does the verb literally express in: “John got the measles from his cousins”? The evident impossibility of specifying such a content shows that whatever content Gricean maxims may help to determine in these cases must be a *literal* content.

²⁰The notions of topicality and semantic openness are easier to develop in terms of the German word “*Spiel*” than in terms of its English translation, “game.” There is no direct English equivalent of “*Spiel*.” The closest English expression is probably the noun “play,” as it occurs in uses like “The children have had ample opportunity for play.” But in German one speaks of “*ein Spiel spielen*,” whereas the constructions “gaming a game” and “playing a play” are (virtually) non-existent. More importantly, “*Spiel*” applies to a broader range of activities than “game.” This comes out even in §66 itself, where Wittgenstein describes ring-a-ring-a-roses and the activity of (solitarily) bouncing a ball off a wall and catching it as “*Spiele*.” The word “game” is not properly applied to these activities. Nor are building sand castles or chasing one another around the yard games—but they are *Spiele*.

²¹Which is not to say that this sentence is not indexical; it is indexical, but only because of the occurrence in it of the tensed “are” (and perhaps also “the cupcakes”—although the contextual determination of contents for this expression seems to be pragmatic, rather than semantic).

²²This tendency is not, however, a universal regularity. “The apples are green” is a sentence of variable content, owing to the presence in it of the polyseme “green.”

²³“Roughly,” since Kittay seems to hold that a term can count as polysemous even if it does not vary in content, as long as its (unvarying) content resists natural non-disjunctive characterization (see [Kittay, 1987, 106–109]).

²⁴As often happens in these discussions, the relation between different authors’ work is easily obscured by variant uses of a common terminology. Kittay describes polysemy as a kind of multiplicity of meaning (see [Kittay, 1987, 113; 150–51]). However, she also distinguishes between “occasion meaning” and “timeless meaning,” the former corresponding more or less to what I have been referring to as “content,” the latter to what I call simply “meaning.” And it is clear that it is occasion meaning that Kittay has in mind in characterizing polysemy as a feature of certain expressions that have “multiple meanings.” By the same token, in another place Kittay suggests that polysemy does *not* entail multiplicity of meanings (see [Kittay, 1987, 107n]); here, however, it is clear (in light of her reference to [Cohen, 1985]) that she has timeless meaning in mind.

²⁵For Kittay’s use of Amelia Bedelia, see [Kittay, 1987, 111]. The test (which takes its name from the famously confused children’s character) works like this: imagine a case in which someone uses a given term to issue an order; now ask: would it be feasible for the person to whom the order was issued to obey it literally, without successfully carrying it out? If the answer is “yes,” then the term in question is polysemous or ambiguous; if “no,” not.

²⁶Support for the claim that Wittgenstein’s notion of family resemblance involves that of topicality comes mainly from passages in which Wittgenstein puts family resemblance to work. Although I do not have space to mount a full defense of this exegetic claim here, I would point to §§156–78 as a place where Wittgenstein tries to demonstrate a philosophical method based largely on the hypothesis that certain terms are topical indexicals. In §162, Wittgenstein imagines a case where a student is to transcribe a printed passage into cursive, by means of a table correlating printed letters with their cursive counterparts. The student can then “show that he is deriving his script from the printed words by consulting the table,” in a variety of ways. The most obvious way would be by writing a cursive “a” for a printed “a,” a cursive “b” for a printed “b,” and so on through “z.” But he would also show that he was deriving the script from the printed text if he always wrote “b” for “a,” “c” for “b,” and so on through writing “a” for “z.” In fact, it “would still be a perfectly good case of derivation by means of the table, even if it were represented by a schema of arrows [correlating printed letters with letters of cursive] without any simple regularity.” (§163) That is: no matter *what* the pupil writes for each printed letter, he still shows that he is *deriving* what he writes from the words on the page, as long as he sticks with a single, consistent method throughout (no matter how arabesque). But what if instead “he does not stick to a *single* method of transcribing, but alters his method according to a simple rule: if he has once written *n* for A, then he writes *o* for the next A, *p* for the next, and so on.—But where is the dividing line between this procedure and a random one?” (§163)

The answer (obviously) is that there is none. But Wittgenstein wants us to draw a lesson from this fact different from the one we are likely to be inclined to draw. The fact that “the meaning [of

“to derive”] seems to disintegrate when we follow it up” (§163)—the fact, that is, that whether we’ve really got a case of *deriving* on our hands becomes unclear as we entertain increasingly less straightforward patterns of transcription—is apt to lead us to the conclusion that even in the cases of straightforward derivation (“a” for “a,” “b” for “b,” etc.), there must have been more to the fact that the pupil was deriving than merely that he consulted the table and wrote down what he did. The problem, as Wittgenstein sees it, is that we are prone to assume that if this “overt” sort of business is sufficient for the pupil’s deriving the written words in the original case, then the same sort of business (of consulting the table, writing down letters according to some method, etc.) must be sufficient for deriving in *all other cases of derivation*. But if we assume that, then of course we shall have to conclude that there is more to deriving than meets the eye (even in the original case); for we can find these overt factors in place even in cases where it is unclear whether we should describe the pupil as deriving. Yet if such factors were sufficient for deriving, there would be no room for such unclarity, and so we conclude that not even with respect to that first case is our claim that the pupil was deriving justified simply by the overt facts of the case. But, although it is easy to lose sight of this, in reaching this conclusion we rely on the unstated assumption that “to derive” must express the same literal content in every use (even if only a “disjunctive” content). For without this assumption we cannot legitimately draw conclusions about what justifies a use of the verb to describe one case from observations relating to what justifies a use of it to describe some (relevantly) distinct case. If “to derive” is a topical indexical having its content responsive (*inter alia*) to topical background facts about those to whom we apply it, then it is perfectly possible that in one (correct, literal, unambiguous) application, the verb expresses a content that could be paraphrased completely in terms of overt behavior, and in another application not. But in that case the unclarity that attends “borderline cases” is not evidence of some “deeper” fact underlying all instances of derivation; it is simply a result of the fact that certain odd cases do not have the sorts of features required to fix a particular literal content for the verb “to derive,” when it is used to characterize them.

This interpretation of §§162–63 in terms of topicality also clarifies the import of §164: what is essential to deriving is not something that is “hidden beneath the surface”—or for that matter lying *at* the surface—of any particular case; for deriving (like reading: see §164) is something for which “in different circumstances we apply different criteria,” in accordance with the different contents that we express by means of the word “derives.” (And the “essence of derivation” cannot be discovered through inspection of any *single* case, since it is essential to a word with the meaning of “derives” that it be capable of expressing *various* contents in relation to various cases.)

²⁷This, of course, is an exaggeration, since not even the most closely scored piece dictates every nuance of play. But for illustrative purposes, I hope the exaggeration may be allowed to pass.

²⁸For more on the rules of conversation and the notion of conversational common ground in particular, see [Green, 1997].

²⁹Again, for more on these points, see [Green, 1997].

³⁰Think of how Socrates sometimes corners an interlocutor so that he has only one option consistent with the Cooperative Principle, namely, to admit that he was mistaken in some previous assertion.

³¹Of course, the use of an expression of *invariable* content can also be closed—recall the case of “protozoan.”

³²The quotations are from [Wittgenstein, 1958, §68].

³³Unless, of course, it is the constitutionality of some such statute that is at issue.

³⁴Or, if we do want to count this as a dead metaphor, then (as L. Jonathan Cohen points out) “most of the language we ordinarily use (including the phrase “dead metaphor”) is dead metaphor.” [Cohen, 1985, 133]

³⁵Here we must bear in mind the distinction between the meaning of a word and what a speaker means by it, or, what it means, on this or that occasion.

³⁶For more on the role of speakers’ intentions in the semantics of common demonstratives and indexicals, see [Kaplan, 1990], [Reimer, 1992], [Perry, 1997], and [Pelczar, 2000].

³⁷In the case of a demonstrative like “that,” the field in which a given precedent is operative may not extend beyond a single conversation, or part of a conversation. With such a term, the slate of precedent is as it were wiped clean at the outset of each verbal exchange, to allow for the construction of a (possibly new) system of precedents in the use of that term (within that exchange).

³⁸Many commentators’ preoccupation with the “criss-crossing” element of Wittgenstein’s development of family resemblance seems to stem from the worry that without *some* kind of internal structure like this, it would be impossible to distinguish a family resemblance term from a useless mark or noise governed only by the rule that “anything goes.” But of course, not *anything* goes in the (correct) literal employment of a semantically open expression, so my interpretation of family resemblance as open topical indexicality does not threaten to precipitate the feared semantic anarchy. The structure we find in the use of a semantically open term does not however rely on some kind of sequential overlapping of the various contents it expresses—any more than the orderliness to be found in the use of a word like “here” or “today” relies on some kind of sequential intersection of the various contents it may be used to express. (On this point I find myself in agreement with a number of commentators, including Hunter (see [Hunter, 1985, 53–68]) and Gert (see [Gert, 1995]).)

On the other hand, if Wittgenstein really *did* intend for this sort of sequential overlap to be understood as essential to family resemblance, he was surprisingly closely aligned with Aristotle and his account of “focal meaning” (see [Owen, 1986]). For instance, the term “healthy” has focal meaning, according to Aristotle, since (1) it expresses a variety of contents in application to various kinds of things (cf. topicality), and, (2) some of the contents it has the potential to express must be defined in terms of others among its potential contents, and all its potential contents (except for the “focal” one) must ultimately be defined in terms of its focal content. Thus “healthy” as it applies to complexions expresses a content that must be defined in terms of the content that the term expresses in application to physical constitutions, relative to which “healthy” expresses its focal content. What makes this content focal is therefore its definitional priority (relative to the other possible contents of the word), and this in turn reflects (for Aristotle, at least) an *explanatory* priority. And what gives one content an explanatory priority over others is the fact that it corresponds to the *essence* of whatever is in question (e.g., of health).

Given my understanding of family resemblance, it would be an exaggeration—but an instructive exaggeration—to describe family resemblance as focal meaning minus the Aristotelian essentialism—and so minus the requirement of overlapping inner-structure whose purpose is simply to subserve such essentialism. (The exaggeration is due to the fact that there is no evidence of an Aristotelian parallel to the notion of semantic openness.)

³⁹It is possible, at least in principle, for an open expression to ossify into a closed one, just as certain conversations can become ritualized over time. (A good example of this is the lengthy formal exchange of greetings practiced by West African speakers of Wolof, at the outset of each conversation.) But although the amount of discretion that speakers may exercise in using an open expression may decrease over time as precedent gets added to precedent, it seems that in most cases—as in common law—speakers who use an open term do not only follow precedents, but occasionally set them. Certainly it seems that precedent in the literal use of the verb “to damage” has not gone so far as to close off the possibility of any further discretion in using it to express literal contents (see §80, §83). This is true even if after a certain point the verb is never used to express any (literal) content that it has not been used to express before; for, as we have seen, it may be at speakers’ discretion to determine *which* of these antecedently expressed contents the verb is to express, in some new kind of application.

⁴⁰I do not mean to suggest that the semantic openness of the verb “to damage” is maximal, or even close to maximal. One cannot, nowadays at least, turn to the literal use of this verb as a major creative outlet. (Certainly a term like “here” or “that” seems to rely more heavily on speaker discretion to determine literal contents for it in various contexts than does “to damage.”) But neither does a typical common law trial afford its jury an enormous amount of discretion, especially compared to the amount of discretion enjoyed by the members of an impromptu vigilante-style trial

(such as that described in *The Ox-bow Incident*); yet the element of discretion in an ordinary common law trial is crucial to our legal system, for all that.

Since conflicting philosophical intuitions are often generated by appeal to a relatively small number of specific cases, a philosophical term would not have to rely for its literal content on speaker discretion more so than does “damage,” in order for the role of speaker discretion in its literal use to be potentially useful in accounting for such “intuition clash.”

⁴¹As we noted at the beginning of this section, there are various reasons why we might find this kind of open endedness useful. These may include a desire to have some multi-purpose linguistic tools at our disposal, together with a reasonable tendency to place more confidence in the abilities of those who actually use a term like “game” or “damage” in particular situations to do so in the most effective ways, than we place in our own ability to foresee what situations might arise, and what it would be best for such a term to express in them.

⁴²However, openness may be important. Consider the term “good.” A corollary of the view that “good” is semantically open is that it can be applied in precedent-setting ways. It may be that basic moral disagreements (such as over scenarios meant to arbitrate between consequentialist and non-consequentialist positions) result when there is disagreement over what precedent should be established, in a given case of application of the term “good” (or, “right”). This would make the resolution of such disagreements on the basis of evidence alone impossible.

⁴³Compare Searle, who simply assumes without argument that a primary or central content of some sort can always be taken to exist. [Searle, 1969, 55]

⁴⁴Compare Wittgenstein’s discussion of reading, and in particular his remark that “in different circumstances we apply different criteria for a person’s reading” (§164).

⁴⁵Thanks to Mitch Green, Harold Langsam, William Lycan, François Recanati, Marga Reimer, Christopher Shields, Charles Travis, and an anonymous referee for *Noûs* for help with earlier versions this paper. And thanks to Kim Haslinger, for introducing me to Wolof.

References

- Baker, G. and Hacker, P. (1980). *Wittgenstein: Understanding and Meaning*. Chicago University Press, Chicago.
- Bambrough, R. (1960). “Universals and Family Resemblances.” *Proceedings of the Aristotelian Society*, 61.
- Braun, D. (1994). “Structured Characters and Complex Demonstratives.” *Philosophical Studies*, 74(2).
- Bréal, M. (1897). *Essai de Sémantique (Science des Significations)*. Librairie Hachette, Paris.
- Clark, W. V. T. (1940). *The Ox-bow Incident*. Random House, New York.
- Cohen, L. J. (1985). “A Problem about Ambiguity in Truth-Theoretical Semantics.” *Analysis*, 45(3).
- Fogelin, R. (1976). *Wittgenstein*. Routledge and Kegan Paul, London.
- Gert, H. (1995). “Family Resemblances and Criteria.” *Synthese*, 105(2).
- Grandy, R. E. (1979). “Universals or Family Resemblances?” *Midwest Studies in Philosophy*, VI.
- Green, M. S. (1997). “Illocutions, Implicata, and What a Conversation Requires.” *Pragmatics and Cognition*, forthcoming.
- Hallett, G. (1977). *A Companion to Wittgenstein’s Philosophical Investigations*. Cornell University Press, Ithaca.
- Hunter, J. (1985). *Understanding Wittgenstein: Studies of Philosophical Investigations*. Edinburgh University Press, Edinburgh.
- Kaplan, D. (1988). “The Logic of Demonstratives.” In Salmon, N. and Soames, S., editors, *Propositions and Attitudes*. Oxford University Press, Oxford.
- Kaplan, D. (1990). “Dthat.” In Yourgrau, P., editor, *Demonstratives*. Oxford University Press, Oxford.
- Kittay, E. F. (1987). *Metaphor: Its Cognitive Force and Linguistic Structure*. Clarendon Press, Oxford.

- Owen, G. (1986). "Logic and Metaphysics in some Earlier Works of Aristotle." In *Logic, Science and Dialectic*. Cornell University Press, Ithaca.
- Pelczar, M. (2000). "Names as Tokens and Names as Tools." *Synthese*, 120(forthcoming).
- Perry, J. (1997). "Indexicals and Demonstratives." In Hale, R. and Wright, C., editors, *Companion to the Philosophy of Language*. Blackwell, Oxford.
- Recanati, F. (1997). "Déstabiliser le Sens." *CREA Rapports et Documents*, 9714.
- Reimer, M. (1992). "Three Views of Demonstrative Reference." *Synthese*, 93(3).
- Searle, J. (1969). *Speech Acts: An Essay in the Philosophy of Language*. Cambridge University Press, Cambridge.
- Sluga, H. and Stern, D., editors (1996). *The Cambridge Companion to Wittgenstein*. Cambridge University Press. Cambridge.
- Strawson, P. (1990). "On Referring." In Ammerman, R., editor, *Classics of Analytic Philosophy*. McGraw Hill, New York.
- Wittgenstein, L. (1958). *Philosophical Investigations*. Oxford University Press, Oxford, third edition. G. E. M. Anscombe and G. H. von Wright (Eds.), G. E. M. Anscombe (Trans.).
- Wittgenstein, L. (1980). *Remarks on the Philosophy of Psychology*. Chicago University Press, Chicago. G. E. M. Anscombe (Trans.).