What I Learned from Al Mackay: Meaning, Synonymy and Translation
A Valedictory Vote of Thanks[*]

Jay L Garfield
Smith College
The University of Melbourne
Central University of Tibetan Studies

1. Memoir and Introduction

First, let me say how surprised and honored I am to be asked to speak on this occasion. First, surprised. I thought that by now we had irrefutable inductive evidence that Al would teach forever. His retirement hence shakes my confidence in the uniformity of nature to its foundations. But I am also honored. For although I only took one course with Al—*Philosophy of Language*—I learned a lot from him, both in that context and in the Faculty Research Seminar during my honors year. And I owe a great deal of my own life in philosophy to Al's teaching, mentoring and example. To be able to participate in this valedictory event is a welcome chance to thank Al for all that he has given not only to me, but to a startling number of my philosophical colleagues.

Let me first say a bit about all of that. Al may have mellowed a bit over the years, but I kind of doubt it, and I hope not. In my days at Oberlin, anyway, Al's philosophy of language course was a rite of passage for majors. It was *hard core.* No nonsense was tolerated. We were expected to come to class not only prepared *in the thin sense*—that is, having read the stuff we were assigned—but to be prepared to *do philosophy, to argue, to fight for truth and clarity.* That demand to be prepared to argue, and to care both about conclusion and reasoning was bracing, and convinced at least some of us that philosophy was meant to *matter.*

[*] Thanks to Guy Newland and to Susan Levin for helpful comments on earlier drafts and for valuable insights regarding translation.

Now of course Al was not the only great teacher and philosopher in the Oberlin department of that day. That was the era of Bob Grimm, Norman Care, David Love, Dan Merrill and Tom Trelogan as well. They were all stars, and each of them inspired many of us. It is fair to single Al out, though, as the one who *scared* us as well. Not only because he demanded so much on paper and in the classroom, but because he showed us that philosophy itself demanded so much, that is a way of life that takes you by the throat and demands that you, in turn, take ideas by the throat—a life that you pursue only at your peril, or perhaps at the peril of your students. That so many of us who studied with Al and his colleagues went on in the profession is testimony to that fact that he—and they—also showed that it is good fun.

Those of us who were lucky enough to be doing honors in the department when Al was writing the Arrow book received a marvelous education in how to do philosophy *hard* in the meetings of the Faculty Research Seminar. And I mean that in several senses: Al worked *hard* to reach clarity, defending his analysis with passion, while listening to and taking on board critique with the greatest intellectual honesty and openness. In doing so, he gave us students an object lesson in taking oneself, one's ideas as well as one's colleagues and *their* ideas *seriously.* And for those of us who were to become teachers, there was another lesson: Take your students seriously, even when they aren't yet really equal players. I have never forgotten that enormous intellectual generosity.

2. The "Brother Argument" and What Was at Issue

But part of being taken seriously by Al Mackay was to find yourself sometimes in the ring with a much more experienced boxer, and to take some hard hits. I want to now revisit a classroom argument I had with Al 35 years ago in Philosophy of Language, resulting in a paper of which he didn't think very much. I still think I was right, and I am sure that he has forgotten the whole thing. Of course I have forgotten most of the details, but returning to the broad outlines of the dispute between us leads to see not only that I was right, but that *he* was, too, for reasons that maybe *he* saw at the time, but which I didn't; reasons that have come to be important to me now.

The paper, I remember, was called, "Let's Get Meaning Out of the Proposition and Back in the Bedroom Where it Belongs" and it was a reply to an argument Al defended in class against the reduction of meaning-statements to synonymy claims. As I say, I've forgotten most of it, but I think I remember the kernel of the dispute between us.

Al asked us to consider the following sequence of sentences:

(1a) "Brother" means *male sibling.*
(2a) "Brother" means the same as "male sibling."
(3a) "Brother" and "male sibling" are synonymous.

(1b) "Brother" means *brother.*
(2b) "Brother" means the same as "brother."
(3c) "Brother" and "brother" are synonymous.

He pointed out that if talk about meaning (sentences of the form "[x] means *y*") reduces to talk about synonymy (sentences of the form "x" and "y" are synonymous), then the substitution sequences (a) and (b) should be equivalent. But they don't seem to be. (a) begins with a contingent and informative sentence, and ends with a contingent and informative sentence. If meaning talk were synonymy talk, it would be sound. (b) begins with a contingent and informative sentence and ends with an uninformative necessary sentence. Hence, Al argued, meaning talk cannot be reduced to synonymy talk.

I argued—if memory serves me right—that the argument is fallacious, and that the fallacy arises from an error about notation, a failure to see what the grammar of statements like (1a) and (1b) is. I argued, I think, along the following lines: Al represented these sentences as *mentioning* the word "brother" in its initial occurrence, the italicized *male sibling* or *brother* was *used,* effectively connecting the term to the meaning. This I regarded as a mistake. No amount of typeface modification, I argued, could crowbar a *meaning* into an expression or into a proposition. All one does is get more language.

SO, I argued, (1a) and (1b) do nothing more than coordinate expressions with other expressions. (1a) is synthetic, and coordinates the expression "brother" with the expression "*male sibling,"* not with some non-linguistic entity magically transported onto

paper or blackboard by typographical mantra; (1b) similarly coordinates the expression "brother" with the expression "*brother."* And the use of italics doesn't do anything; so, (1b) is *not* contingent or informative. Argument (a) does take us from an informative premise to an informative conclusion, while (b) takes us from a boring analytic premise to a boring analytic conclusion, and the master argument, which relies on the claim that both (1a) and (1b) are informative while (3a) and (3c) differ in this dimension, fails. Now, I am sure that I did not argue that meaning talk *is* synonymy talk. I didn't care much about that. But I did argue that changing typeface only amounted to a different kind of quotation, and that this form of argument couldn't show what Al wanted it to show. Needless to say, by the way, Al would have none of this.

Why dredge all of this up now? Well, first of all, because I still think I was right about the argument, I have a hard time letting things drop, and this is my last chance to win this fight. But also because I have also come to see that *Al* was right, not about the argument, but about the conclusion. And my own philosophical life has never really strayed too far from the philosophy of language. When I heard that Al was retiring this argument was the first thing that came back to my mind. So, I thought I'd have it out once and for all, and concede what needs to be conceded.

3. Dot Quotation and Vindication

Well, here's the first thing to say. Little did I know then that I was fumbling badly for the insight to which Sellars had come decades earlier, and that was going to drive my own work in epistemology and the philosophy of mind for the next few decades, the idea that the predicate *means* is really a form of the copula, and that the expressions that occur to it right are in fact in a curious kind of quotational context, the dot-quotational context. On the Sellarsian view, Whenever we assert something of the form *X means Y,* we are actually classifying '*X'*, saying that it plays the same role in its language that '*Y'*plays in its language. This is now familiar to all of us. But when I left Oberlin for Pittsburgh, it hit me like a thunderbolt.

And of course I saw this as vindication in my battle with Al MacKay. Now I could make the point precise:

(1a) “Brother” means *male sibling.*
(2a) “Brother” is a •male sibling•.”
(3a) “Brother” and “male sibling” of the same type.

(1b) “Brother” means *brother.*
(2b) “Brother” is a •brother•.
(2c) “Brother” and “brother” are of the same type.

When we see these two triples of sentences this way, the aura of difference disappears. In each case all we have are statements about functional classification. One set contains informative sentences; the other does not. We have no argument against the replacement of meaning talk with synonymy talk. But my victory turns out to be Pyrrhic. That is, while I am right that the argument fails, its conclusion is nonetheless true, for reasons which I was only to appreciate much later in life, and, ironically, for reasons that have everything to do with seeing meaning as functional classification.

4. Life as a Translator

As it turns out, philosophical careers take one in unexpected directions. Al, for instance, has drifted from the philosophy of language and the philosophy of economics into Ancient Greek philosophy of late. For me, the drift took me from logic and cognitive science to Indo-Tibetan Buddhist philosophy and into the life of a philosophical translator. The practice of translation, moreover, took me as a philosopher into reflection on the practice of cross-cultural interpretation and the nature of translation. That is, I have found myself thinking a lot about cross-cultural hermeneutics. And talk about meaning is not as simple in that domain as one might hope.

Translation turns out not to be what I once thought it was, *viz.,* the replacement of an expression in the source language with an expression in the target language that means the same thing. Those of us raised on Davidson and Sellars, or even on Quine, are in for a surprise when the rubber of real texts hits the road of real translation. It turns out that we *can* both recognize two sign systems to be *languages*, and *understand* them both, and find it impossible to find expressions on one language that are equivalent to expressions in the other. Davidson was just wrong about that.

First, a few words about translation and its goals, in a very informal sense. When we translate, we produce a new text—one in a *target* language—that bears a complex relation to another text, the original text (usually—but sometimes even a translation itself) in the *source* language. Kemp Smith and Guyer, for instance, produced distinct English (target) texts called *The Critique of Pure Reason,* each of which bears the translation relation to the German *Kritik der Reinen Vernuft.* A translator juggles and weighs a number of desiderata, which, as Dorit Bar-On and Luis Goméz, in their respective masterful essays on the craft, point out, are generally in tension with one another. One hopes to produce a text whose constituent sentences have the same truth-conditions as those in the source language; one hopes to convey the author's literary style; one may hope to preserve such aesthetic features as rhyme, meter or lexical allusion. One might wish to achieve lexical consistency, with a close to one to one mapping between source and target lexemes; one might want to preserve grammatical constructions, and implicatures, etc… But most of all, one wants the target language text—the translation—to have roughly the same literary form, and to be as readable as, the source text. A text in straightforward colloquial prose in French should not emerge in highly technical English encumbered by explanatory notes, bracketed material, etc… Nor should a highly technical, precisely written text emerge as an informal, colloquial discussion, however more pleasant that might be to read. As I said, these demands always conflict, and translation is always a matter of failing in the least egregious way possible, never of succeeding.

Let me illustrate by example, both examples that vex me. First, let's try moving from Sanskrit to English, using a common philosophical technical term that has even made it into English in popular culture, *dharma.* The root of the term is *dhr,* which has the sense *to hold, to support.* In Sanskrit, *dharma* can denote a *doctrine* (something that one holds on to, or as we would put it, with similar etymology, something one *maintains);* it can denote an *entity* (something that *holds* together); or a *property* (something *held*); *virtue* (something one *ought to hold on to,* or that *holds us on the straight and narrow*); the most fundamental constituents of reality (that which *holds together* to constitute things).

It is important here to emphasize, that if you are thinking in Sanskrit (or in Tibetan, which has coined a term *chos* has the same semantic range, with none of the etymology)

this is not a case of homonymy, like *bank* and *bank* in English. *Dharma* is not ambiguous, but is simply a term with a broad semantic range, like the English word *game.* Thinking in Sanskrit, *dharma (virtue)* is as similar to *dharma* (*entity)* and *dharma* (*property)* as a game (Football) is to a game (monopoly) or a game (Sudoku). Of course *dharma* is a technical term, and *game* pretty non-technical (though see game theory), but in this respect they are similar: they each have broad semantic ranges, but are not therefore *ambiguous.* So, how do you translate *dharma* into English? Easy. You find the word that has the same meaning—you know—the word that means *entity, property, virtue, doctrine, etc…*

I have to do this all the time when translating Buddhist philosophy from Tibetan into English. No matter what Davidson told us about the necessary intertranslatability of languages, no term is forthcoming, only circumlocutions or extended footnotes, which are different in kind from translations. You don't translate a text by writing a commentary on it; you comment on a text by doing that. It makes sense to say of a set of explanatory notes that they do not constitute a translation, but an exposition, in virtue of the different authorial voice, the different literary form, and the fact that what is literally stated in the comments is never stated, and does not need to be stated in the source text.

Or one could leave the term untranslated. This does have some virtues: As students in a seminar of mine commented with regard to the difficult Sanskrit term *pramāṇa* (*epistemic instrument, warrant, warranted cognition, evidence, etc…)* the English options are each so bad that it is best to teach them the meaning of the Sanskrit term and its theoretical background term and just use it. That is often how loan words work. So from Buddhist texts themselves we acquire the English terms *nirvana, Buddha, karma, lama, guru, mandala, etc…* But there are problems here, too, of an especially insidious kind. For once these terms migrate into a target language like English, they develop their own new semantic range, and no longer accurately translate the source terms with which they are homophonous and from which they derive. Their homophony then becomes a trap—creating an illusion of precise lexical fidelity where there is none.

Or one might consider translating the same term with different non-cognate, non-synonymous English terms on different occasions, even though in the original it conveys

a single meaning. This, despite being the best option in many cases, is in an important sense, a mistranslation. We fractionate meaning that was unitary in the source, and where that unitary broad range played a critical role, into a set of unrelated meanings in the target where the very point about the breadth of range is lost, or available, again, only in notes.

Here's the other example, in the other direction, one that bedeviled me when translating Western philosophy into Tibetan. We were writing about and translating texts by Schopenhauer, and so one cannot avoid the term *will.* Now, having grown up in the West, we take the will pretty much for granted, worrying about whether it is free, strong, weak, etc… and assuming that because we talk and think so freely about it, there is obviously something about which we so freely talk and think. Think again. And let's talk. Look inside. Do you find a will? Have our best neuroscientists found the will somewhere in our brains? "Will" is neither a term of simple observation nor a theoretical term of any developed science. It is instead a term of philosophical art that has migrated into popular culture in a quasi-Jonesian style.

Jones, however, in this case was not a mythical character dreamed up by Sellars. He was an actual saint—Augustine—who invented the will to solve a very particular problem, the problem of theodicy raised by Adam's fall. If you are bothered a lot by Adam's fall, and whether an omnipotent, omniscient and omnibenevolent God was responsible, you'd better have a will at your disposal, and its freedom from the causal order as an accessory, or your theology becomes incoherent. But suppose that you are no bothered by that story or that theology, perhaps because you never heard of it. Then would you have invented this *tertium quid* to mediate between intention and action, free from the causal order, and beset by innumerable philosophical inconsistencies? And would you have decided a millennium and a half later that it is the *ding an sich?*

The answer is *no.* And so we find that in Sanskrit and Tibetan there is no term that translates the English/German term *will.* And there is no debate about its freedom in that literature, etc… So, in translating Schopenhauer into Tibetan, we had to seek the term in Tibetan that means *that thing which might be uncaused that mediates between intention and action.* You might as well try to translate *Santa Claus,* looking for that term that

denotes a fat guy who flies through the skies at night pulled by twelve flying reindeer. There just is none. We made one up, with a long explanatory footnote. That's not really translation, either. That is the creation of a calque.

Things can get interesting syntactically, too. Many of the world's languages have evidential markers, terms that convey (though this is a complex story) the source of evidence a speaker has for the claim she is making, such as hearsay, inference, or direct perception. We convey that kind of information in English using propositional attitude verbs or epistemic modals, such as *I infer that,* or *I saw that* or *It must be that.* But those create embedded contexts, and change the truth-conditions of the matrix sentence. If I say *I infer from perceptible evidence that it is raining*, and I don't, even though it is raining, I have said something false. But if I report using a specific-inference evidential that it is raining, I have said something true (though pragmatically infelicitous) even if I don't rely on that kind of evidence. So, translating between evidential and non-evidential languages is impossible. What is illocutionary in one must be locutionary in the other. One resorts to circumlocution and commentary, not real translation. (Dorit Bar-On points out that gender generates similar syntactic bars to translation.)

5. Too thin a Diet of Examples

So how did so many people decide that translation has to be in principle possible, that we can always find a cross-linguistic synonym or close enough, that there is always something to put between the dot-quotes? The answer is the reliance on what Wittgenstein felicitously called "too thin a diet of examples." We focus on common nouns and predicates,: *rabbit*, *red*, *it is raining,* and we find that we can find really good equivalences for these, plus or minus a bit. We then treat them like the frictionless planes of physics, and decide that the translation of more problematic terms is just like those, only a bit harder. We convince ourselves that the only way to recognize that something is a language is to have a complete translation manual for it, despite the fact that we recognize *many* languages, and have *complete* manuals for *none*, and go on our merry way. In short, we ignore as theoreticians of translation (or, more accurately, as philosophers of language who never notice that translation is where the rubber of semantic theory hits the road of actual semantic problems) the difference between what

we might follow Williams in calling *thick* and *thin* terms. Thin terms are easy. Thick ones are often really hard. That's why there is no good German term for *schadenfreude.*

One way of seeing why this is so is to reflect on what language really does. Most of us, myself included, have abandoned the view I defended in my Oberlin honors thesis all those years ago—that language is a tool for *picturing* reality.

Most of us can now quickly rehearse a range of arguments for the falsity of that position. And most of us have drunk enough Wittgensteinian Kool-Ade to think about language as interwoven with something called "forms of life." But often, when we then turn to the task of natural language semantics, we fall back into our neo-Fregean habits of looking for the denotations or intensions of terms as specified by evaluation functions that map terms neatly to individuals, sets or functions. And that encourages us to think that all we need to do in translation is to discover trans-linguistic synonyms, words mapped by evaluation functions from distinct lexical domains to the same semantic values. Since without such functions, we think, there would be no meaning, and since there is meaning, there must be such functions, and we are off and running.

But that really is to think of meaning in the wrong way, driven by what seems to work so well for the thin terms that constitute not the thin end of useful theoretical wedge, but the visible tip of a perilous linguistic iceberg. To the extent that we really appreciate the consequences of the late Wittgensteinian revolution, and take seriously the idea that language is a social coordination and communicative tool intimately woven into our entire way of being in the world, we have to take equally seriously the impossibility of completing any version of the project of denotational semantics for natural language, and the more general impossibility of detaching a theory of meaning from an account of ideology (think Augustine or Buddhism) or pragmatics (think evidentiality). Language, in short, is one of the many things we *do,* not a mirror we hold up to the theatre of non-linguistic action. Translation then is not the selection of semantic equivalents, but a complex kind of approximate re-enactment, whose success or failure is determined, just as that of any performance, by a rough consideration of the compromises made in staging. And there always are compromises. That is why translation is such a dismal business.

6. Why Meaning is Not Synonymy

All of this is a roundabout way of saying that I finally realized late in my career that Al was right all along. Talk about meaning is not talk about synonymy after all. But not because of the brother argument. Al was right because there is no such thing as synonymy, even though there has to be meaning. To be sure, some—most notably Quine—have argued against the reality of meaning altogether. But any such argument is in the end self-refuting. If there is no meaning, then there is no way to use language to say that. That point was urged by the Nayaiyikas two millennia ago in India, and was recognized by Wittgenstein in the last century in Europe. On the other hand, it does not follow from the truism that there is meaningful language that there are *meanings* to which words are related, or that can be named by terms in funny typefaces in sentences or that any terms bear a strict synonymy relation to one another. That point was recognized by Nāgārjuna two millennia ago, and also by Wittgenstein, as well as by Sellars, in the last century.

How does one preserve meaning while jettisoning meanings and synonymy? Easy. One adopts a use theory of meaning according to which, as Wittgenstein put it, "the meaning of a word is its use in language." One can then translate *x means y* into *x is a •y•* in the familiar Sellarsian fashion, creating a name for the role that *y* plays in it language. But here we name a role, a use, and hence a property of a term, not an abstract entity to which it is *related.* The apparently real meaning *relation* is jettisoned, and with it the second relatum to which it is supposed to relate words.

Classification, of course, is typically both vague and interest relative, including the classification of roles. We might, for some purposes say truly that the President plays the role in the United States that the Queen does in Australia (Head of State); or truly *deny* that claim and say that the President instead plays the role of the Prime Minister (Head of Government). Or say that *nothing* plays the role in America that either does in Australia (having one of those functions but not the other). Which role-classificatory claim is true depends on one's interests, since identity of kind depends upon which kinds count, and which kinds count depend on what cares and concerns we bring to the table. And then there is vagueness. Do coaches play the same role in professional baseball that they play in little league? Yes, and no. So, if we shift from talk about relations to meanings

captured by characteristic functions to classification of roles, we lose a lot of precision in the process of gaining a lot more insight.

And there goes synonymy as well. For synonymy, if it is not to be the higher-order relation of bearing the meaning relation to the same meaning—a non-starter in the present framework—it can only amount to the relation of having exactly the same use. But no two words have exactly the same use. In the end, the close observation of language yields only tighter or looser similarity relations. Just as so-called identical twins are identical only in certain respects, and when considered *tout court* only very similar to one another, even the words closest in meaning to one another are only very similar—interchangeable in some contexts, but never in all. And of course there is no great loss here. The creation of sortals and concepts of all kinds proceeds on the basis of similarity relations and family resemblances; there is no special problem for the construction of metalinguistic sortals.

So, if there is meaning, and there is no synonymy, it had better not be the case that meaning reduces to synonymy, or that talk about meaning can be replaced with talk about synonymy. I certainly did not see that 35 years ago, but I'll bet that Al did.

7. Heterotranslation and Autotranslation: the Universality of the Hermeneutic Predicament

For those like me, who devote substantial portions of our lives to the thankless task of translation, and for all of us who are consumers of translations—and I suspect strongly that that includes everyone in this room—such an analysis raises interesting hermeneutical questions about the nature of translation itself. It is tempting to think that when we translate we aim to find words that bear the same meaning in the target language that those we are translating bear in the source language—that we are on a quest for cross-linguistic synonymy, if not for words, than at least for phrases, or sentences. Otherwise, one suspects, we run the risk of composing not a translation, but a paraphrase.

But this can't be right. We are on the one hand not just trying to paraphrase. Something counts as getting it closer to right (though, as one learns, nothing counts as getting it *right*—this is source of the despair of all who practice this craft, and the source of the frustration for all who rely on that craft). Here's what I think: We aim to come as close as

we can to satisfying the following counterfactual: We try to say *in* the target language what the author of the source language *would have said* had *she* composed the text in the target language. That is, we aim for the tools she would have picked up had she had our toolkit in front of her instead of her own.

Note that meaning drops out of the picture here. Translation, I have come to believe, is more like theatre than it is like lexicography: we aim, as does an actor, to perform speech acts that someone we never knew would have performed had she found herself in a situation she could never have imagined. No wonder it is so hard, and no wonder different translators translate so differently. It is for the same reason that different actors play Hamlet so differently. It is for this reason as well, that just as we value different ways of performing the melancholy Dane, we value different translations of the same text. This is not simply a phenomenological report, although it certainly is that; it is also the only sensible account of what translation *could* be if my reflections on meaning are correct. If meaning really is use, and if words really are tools, and if none of them are identical in use to any others, translation can be nothing else.

Why reflect on translation, though? Of what interest could these reflections be to anyone smart enough not to translate? Well, here's what I think: *Translation* practice tells us a lot about *linguistic* practice more generally, and so a lot about *language*. We can put the point this way: what goes for heterotranslation (translation *between* languages) goes for autotranslation (translation within one language). While we might not notice it, we autotranslate all the time. We do so when we employ indirect discourse. We do so when we reflect on our own thoughts and those of others. We do so even when we attribute propositional attitudes to ourselves or to others, and given our social nature and the ubiquity of that activity in our lives, that is a lot of translation.

Translation, then, whether within or between languages, is coextensive with interpretation, and so involves this odd counterfactual projection. Projecting counterfactuals with regard to expression is a notoriously open-ended, interest relative and context sensitive affair—the hermeneutic predicament in which almost anything can turn out to be relevant, in which what we find in the foreground—that which is the immediate object of interpretation—often has significance only in relation to a

background that sets the possibilities and the presuppositions of interpretation, and itself consists of semantically evaluable phenomena, phenomena which themselves demand interpretation against yet more distant horizons.

This brings us back to lessons taught by those other giants of the 20th Century, Heidegger and Gadamer, about the universality of the hermeneutic predicament. All meaning emerges only in acts of interpretation, and all interpretation, all understanding, emerges from *applicatio,* which always involves our taking up a message and bringing it into the purview of our own cognitive horizon. The meaning we take to reside in words or other signs is only the anticipation of the semantically constitutive interaction with a reader, a listener or a user. And since the distinct presuppositions, purposes and backgrounds that condition each interaction—each use of a word or a text—is always shifting, *meaning* is always shifting, always unstable, always a work in progress. Understanding is an ongoing *activity*, not an accomplishment.

Hermeneutic activity, moreover, is endlessly circular. Just as the meaning of any term on any occasion of use depends upon the entire context of its use—use both by the author and by the interpreter—the symbolic context that gives meaning to each part depends for its significance on the meanings assigned to those constituents. And just as we are constituted as readers, authors, speakers and listeners by the prejudices (fore-judgments), projections, and questions we bring to the texts we produce and interpret, those texts affect us, altering those projections, prejudices and questions so that when we next encounter those texts, new readers are reading new texts. Given these moving targets, it is no wonder that talk of synonymy in the real world of linguistic praxis is a bit of a joke.

But it is also important to understand that linguistic practice is always embedded in and inextricable from extralinguistic practice. It is easy for philosophers, especially philosophers of language, to lose sight of this fact, though once one turns one's attention to it, it is a bit of a truism. *Homo sapiens* we may be, but our sapient activity is hardly *sui generis.* At bottom, as Aristotle noticed, we are *social animals.* Our thought and our speech are social coordination devices, biotechnologies for efficient commerce with our conspecifics and the environments we create.

A further hermeneutical circle is hence constituted by the reciprocal dependence of our symbolic activity on our entire range of individual and social purposes, and the dependence of those purposes on our thought and speech. Just as the world we inhabit is a *lebenswelt* constituted by interpreted objects, events and colleagues, determined in its character in part by our interpretations, our future interpretations are conditioned by our life in that world. Meaning is hence not only labile, but also porous. Its determinants extend well beyond language itself; at the same time our symbolic activity projects meaning into the non-symbolic world, always extending the possibilities of interpretation. Anyway, that's where the philosophy of language has taken me, once I started to really worry about language.

8. What I really learned from Al Mackay

I learned a lot from Al MacKay. One lesson is that Philosophy is hard, and insight is hard to come by. If a problem is worthy of attention, a whole lot of smart people have already attended to it, and the low-hanging fruit is long gone. Philosophy therefore can only be pursued with vigor and commitment.

I also learned, both in the classroom and in the research seminar that you only get anywhere through argument and interchange. A lot of people think of philosophy as a solitary activity, best pursued by the lone scholar sealed in a cave with only a few books and a laptop. Any of us who have been doing it for a while know how false that is, how much we depend upon our colleagues and interlocutors. But you have to learn that, and I am immensely grateful that I was taught that lesson early, and that after Oberlin I always sought to do philosophy as an ensemble practice.

I also learned from Al the valuable lesson that when you do philosophy you are mostly wrong. Al showed us through his powerful critical approach to the subject that interesting ideas are almost always false, or at least half-baked and susceptible of refutation. But this, we learned, is no reason to be depressed. The bar for philosophical insight is not set as high as truth. Sometimes it's enough just to be interesting—to move the discussion along.

And there is no shame in being wrong. Al taught me as well that you learn more from being wrong than from being right. Philosophy and philosophers each develop through what Imre Lakatoś felicitously called "a process of conjecture and refutation." We often learn our most important lessons when we see just why our best ideas turned out to be dead ends. Commitment to error is never time wasted, only fuel for growth.

But I suppose that the most valuable lesson Al taught me was the one that has defined my life. He taught me this by example: there is nothing more fun than teaching philosophy. And Al has had a lot of fun. So have we, his students, and we thank him for it.