

Why Philosophers Shouldn't Do Semantics

Herman Cappelen

"All crafts, trades and arts have profited from the division of labour; for when each worker sticks to one particular kind of work that needs to be handled differently from all the others, he can do it better and more easily than when one person does everything. Where work is not thus differentiated and divided, where everyone is a jack-of-all-trades, the crafts remain at an utterly primitive level" I. Kant

1. Introduction

The linguistic turn provided philosophers with a range of reasons for engaging in careful investigation into the nature and structure of languages. However, the linguistic turn is dead. The arguments for it have been abandoned (for good reasons). This raises the question: why should philosophers take an interest in the minutiae of natural language semantics? I'll argue that there isn't much of a reason – philosophy of language has lost its way. Then I provide a suggestion for how it can find its way again.

Two pathways into the topic of this paper:

1. Currently philosophy journals publish papers on issues that can be described as contributions to *natural language semantics*. Paradigms include work on the semantics for: anaphora, indefinite and definite descriptions, quantifiers, epistemic modals, conditionals, and moral language. Question: for those who have given up the linguistic turn in philosophy (more on what that involves below), isn't working on these topics a job for linguists? Linguists are trained in phonology, syntax, semantics, pragmatics, and a range of related fields. That provides a much better preparation for doing work in

semantics than spending years in graduate school writing papers on a wide range of philosophical topics unrelated to natural language semantics. Much of linguistics – and semantics in particular – has its historical origin in work done by philosophers. However, when philosophy spawns a new discipline, philosophers stop doing detailed work on the topic of that discipline. For example:

- a. At some point what we now call ‘psychology’ was done by people who called themselves ‘philosophers,’ but now psychology is its own discipline. At this stage, philosophers don’t contribute to, say, diagnostic criteria for PTSD. It would be bizarre for the *The Journal of Philosophy* to publish a paper that tried to contribute to that literature and it would be an unfortunate use of academic resources to have someone without training in psychology have that specialty.
- b. Similarly, what we now think of as economics was once done by people who described themselves as ‘philosophers’. But now economics is its own discipline. *Mind* wouldn’t and shouldn’t publish papers on, say, inflation rates in Argentina. It would be an unfortunate use of academic resources to have someone without training in economics to have that as a specialty.

Why, then, are philosophers still contributing to natural language semantics? Isn’t that just like philosophers trying to contribute to psychology or economics?

2. Here is a second way into the topic of this paper: The first few chapters of Timothy Williamson’s *The Philosophy of Philosophy* is a sustained attack on the linguistic turn in philosophy. For the most part that attack is successful. For those, like me, who are convinced by Williamson’s arguments, that raises the question: What follows about the philosophical significance of philosophy of language – and natural language semantics in particular ¹ - if we take on board Williamson’s lessons? At one point in his book Williamson says: “Although philosophers have more reason than physicists to consider matters of language or thought, philosophy is in no deep sense a linguistic or conceptual inquiry, any more than physics is.” (2007: 21) Do we really have more of a reason than the physicists to take an interest in questions about language? Williamson has an answer to that question. I think it fails. In the second half of this paper, I provide what I think is a better, though still tentative, answer.

Here’s a preview of the view I will end up with: I will distinguish between the semantics of natural languages and philosophy of language. Philosophy of language is not (and arguably doesn’t even include) natural language semantics. Philosophy of language, properly construed, has two parts:

¹ More on the difference between natural language semantics and philosophy of language below. For reasons of simplicity, the focus in this paper is on the role of semantics in philosophy. Much (if not all) of what I say applies directly to most work in pragmatics. It shouldn’t be too hard for the reader to apply at least many of the lessons from this paper to work on, say, scalar implicatures and presupposition projection.

First, it includes work on the foundations of linguistics (as philosophy of physics and mathematics work on the foundations of those disciplines – more on the notion of ‘foundation’ below).

Second, it includes work on *conceptual engineering*. That was how philosophy of language started and should be seen as its proper domain. From this point of view, the move towards a focus on natural language semantics has been an unfortunate mistake that has contributed to the marginalization of the philosophy of language within philosophy.

The proposal in the last part of this paper – that philosophy of language should in large part be focused on the theory of conceptual engineering – makes philosophy of language again central to all areas of philosophy (and beyond). So one way to think of this paper is as a new argument for why philosophy of language should be construed as a form of first philosophy (though, ‘first philosophy’ has all kinds of misleading hierarchical associations, so better to think of it this way: it’s a subdiscipline of philosophy with very immediate relevance to all other philosophical subdisciplines (and beyond)).

2. The Abandonment of the Linguistic Turn

In the introduction to *The Linguistic Turn*, Rorty says:

The purpose of the present volume is to provide materials for reflection on the most recent philosophical revolution, that of linguistic philosophy. I shall mean by “linguistic philosophy” the view that philosophical problems are problems which may be solved (or dissolved) either by reforming language, or by understanding more about the language we presently use. (1967: 3)

Timothy Williamson, in *The Philosophy of Philosophy* says, ““The linguistic turn” has subsequently become the standard vague phrase for a diffuse event – some regard it as the event – in twentieth-century philosophy, one not confined to signed-up linguistic philosophers in Rorty’s sense. For those who took the turn, language was somehow the central theme of philosophy.” (2007:10)

Succinct statements of this kind of view can found in both Ayer and Dummett:

[T]he philosopher, as an analyst, is not directly concerned with the physical properties of things. He is concerned only with the way in which we speak about them. In other words, the propositions of philosophy are not factual, but linguistic in character – that is, they do not describe the behaviour of physical, or even mental, objects; they express definitions, or the formal consequences of definitions. (Ayer 1936: 61-2)

Only with Frege was the proper object of philosophy finally established: namely, first, that the goal of philosophy is the analysis of the structure of thought; secondly, that the study of thought is to be sharply distinguished from the study of the psychological

process of thinking; and, finally, that the only proper method for analysing thought consists in the analysis of language. [...] [T]he acceptance of these three tenets is common to the entire analytical school. (Dummett 1978: 458)

Williamson goes on to make a sociological/historical point: It's empirically false that philosophers working in the analytic tradition assume that language has such a central role. That view is not in the common ground among philosophers in the tradition we loosely describe as 'analytic'. Metaphysics provides a good illustration:

Much contemporary metaphysics is not primarily concerned with thought or language at all. Its goal is to discover what fundamental kinds of things there are and what properties and relations they have, not to study the structure of our thought about them – perhaps we have no thought about them until it is initiated by metaphysicians. Contemporary metaphysics studies substances and essences, universals and particulars, space and time, possibility and necessity. Although nominalist or conceptualist reductions of all these matters have been attempted, such theories have no methodological priority and generally turn out to do scant justice to what they attempt to reduce.

The usual stories about the history of twentieth-century philosophy fail to fit much of the liveliest, exactest, and most creative achievements of the final third of that century: the revival of metaphysical theorizing, realist in spirit, often speculative, sometimes commonsensical, associated with Saul Kripke, David Lewis, Kit Fine, Peter van Inwagen, David Armstrong and many others. (Williamson 2007: 18-19)

And this point generalizes: it's simply an empirical fact that most of the leading contributors to philosophy of mind, logic, mathematics, epistemology, ethics, perception, and so on don't see themselves as working in philosophy of language or as engaged in an effort to analyze language. Understood as a description of a shared assumption in that tradition, the claim is empirically false.

An ardent defender of the linguistic turn might reply:

Defense of Linguistic Turn: The core tenets of the linguistic turn are true even if their truth is not part of the common ground among analytic philosophers (or philosophers more generally). Philosophy, even when it's not obviously about language, is about language in a disguised way. The philosophers who don't realize this are confused.

Alternatively, a proponent of the linguistic turn can insist that the best (or only) way to solve philosophical questions that are not about language is to study language (and those who don't proceed in that way are making a mistake).

Here's an important juncture in this paper: I'm going to assume that these kinds of defenses fail and I'm not going to argue it. To do so could take up the rest of this paper, but it's not what I

want to focus on.² Instead, I want to pick up the thread of the discussion at the point where Defense is rejected. In other words, for the rest of this paper, I will make the following assumption:

No Priority for Philosophy of Language (No Priority): It is not the case that questions about justice or race or grounds or consciousness or naturalism or God or social ontology or explanation or causation or knowledge or beauty or action or validity or attention or perception or (continue with more or less any question not directly within philosophy of language) are about or are best solved by studying the properties of words or sentences.

Assuming No Priority, I then ask: how should we think of the role of philosophy of language in philosophy? In particular, what's the relevance of natural language semantics to the rest of philosophy and why shouldn't natural language semantics be handed over to linguists?

3. What happens to philosophy of language when the linguistic turn is rejected?

The first part of Williamson's *The Philosophy of Philosophy* is a powerful and successful attack on the linguistic turn. You might then expect Williamson to conclude with a very negative view of philosophy of language. But that's not what happens in the book. On the contrary, he argues that reflection on the structure of natural languages has an important and prominent role in philosophy. Before I consider Williamson's argument for that, I first consider three answers to the question "Why should those who reject the linguistic turn care about the properties of natural languages?" These are:

- (i) Abandonment
- (ii) Accidental Historical Connection
- (iii) Piecemeal and Topic Specific Significance

I then, in section 4, turn to Williamson's own answer. I reject that answer too, and in section 5, I outline my own account of why philosophy of language (but not natural language semantics) should be a central part of philosophy.

3.1. Abandonment

This is the most pessimistic reply. Those advocating it hold that there's no good reason why philosophers should take an interest in the nature and structure of languages. Here are three arguments for Abandonment:

² I think the central arguments against that kind of view can be found in Williamson 2007, and I add some arguments in my 2013, chapter 10.

(i) The Linguistic Turn provided the **only** good reasons for thinking that philosophy of language was a significant part of philosophy. According to proponents of the linguistic turn, all philosophical questions are either **best** answered by going metalinguistic and asking questions about language, or are **only** answered in this way. Accompanying this view is, as we have seen above, the claim that questions in philosophy that don't *appear* to be about language really *are* about language. Proponents of Abandonment then say: The linguistic turn was completely wrong and, once this is realised, there's no motivation whatsoever left for thinking that a focus on the structure of natural languages is philosophically significant. The arguments given by proponents of the linguistic turn were the **only** strong arguments for not just shipping the study of natural language semantics over to the linguists.

(ii) Suppose the first argument doesn't convince you: you think that there are some reasons why issues about natural language semantics are more significant for philosophy than for, say, physics, law, history, or psychiatry. The second argument for Abandonment appeals to the division of labour between disciplines:

The Division of Labor Argument: In order to understand natural language semantics you need the kind of training that linguists get. You need training in phonology, syntax, semantics, pragmatics, and forms of morphology, and all the other topics that you get a training in when you do a graduate program in linguistics. Philosophers don't spend their time being trained in this way and so it's an unwise use of academic resources to have philosophers work on these issues.

According to this argument, having philosophers work on figuring out the semantics for a natural language is like having philosophers contribute to the literature on cancer treatment – we just don't have the training and it would be silly to have philosophers do that even if some of us have some half-baked ideas about how to do it (and took some summer courses on the side).

(iii) Here is a closely related argument – an argument by analogy:

The Argument by Analogy: very few philosophers are of the view that phonology papers should be published in philosophy journals. Very few philosophers are of the view that papers on the syntax of, say, Romanian, should be published in a philosophy journal. Suppose you agree that this would be weird – so phonology and Romanian syntax is out. Question: why should papers on the semantics of English definites or indefinites be any different? Proponents of Abandonment say: There's no relevant difference. Since we don't publish papers on phonology or Romanian syntax, we should not publish papers on how 'a' or 'the' work in English. It is equally bizarre. It's also equally parochial: some languages (such as Latin and Polish) don't have articles, and in some that do, they behave differently from articles in English. So while it's not uninteresting to understand how these expressions work in English, it is very much like understanding some idiosyncratic

parts of Romanian syntax or phonology - it's interesting, but why would philosophers be doing it?³

We don't notice its weirdness because we're still stuck with prejudices from the linguistic turn. We still see philosophical significance through the lens of the linguistic turn, but this is clouded vision. Once we fully take on how wrong the linguistic turn was, then philosophy journals should stop publishing papers on natural language semantics.

I think philosophers (and philosophers of language in particular) have failed to take these challenges seriously enough. They are hard to respond to. At the end I will sketch a reply. The reply relies on drawing a distinction between philosophy of language and natural language semantics, suggesting that while we should leave natural language semantics to the linguist, philosophy of language should either focus on foundational issues or be a form of what I call *conceptual engineering* (which is something we're better equipped to do than linguists are).

3.2. Accidental historical connection: philosophy is just whatever people in philosophy departments do

I turn now to the second reply to the question "Why should those who reject the linguistic turn care about the properties of natural languages?" I call it *Accidental Historical Connection*. To see what I have in mind here, consider first an answer to the completely general question of why something is (or should be) classified as philosophy (the question addressed in this paper is a more specific version of that very general question: why classify natural language semantics as philosophy). In my view, it's hard to do much better than appeal to an institutional-historical answer: *Philosophy is whatever has happened to be classified as 'philosophy' for various contingent and weird historical reasons*. In particular, there's no thematic unity and no methodological unity. As a rough, generic, answer: *If you want to know what philosophy is, go look at what people hired in philosophy departments work on*. That's what philosophy is⁴.

If, like me, you're attracted to that kind of view, then natural language semantics belongs to philosophy because there's a bunch of people hired in philosophy departments who work on these topics. The deeper question of whether this is 'really' philosophy assumes some kind of essence to philosophy, but there is no such essence.

According to this view, issues having to do with natural language semantics are more significant to philosophy than, say, to law or history or archeology because – for weird historical reasons – there's a significant and influential group of people hired by philosophy departments who work on those issues, but that's not true in law, history or archeology departments. On this view, natural language semantics has no more intrinsic - e.g., thematic or methodological - connection

⁴ When worked out in a more detail, this becomes the philosophical analogue of what's sometimes called 'the institutional theory of art' - of course the two theories differ in what histories and institutions are appealed to.

to philosophy than to law, history, etc. It's just a historical accident that the people working on those issues are on the same floor as those working on ethics, metaphysics, and so on. But those accidental historical connections are what make for an academic discipline.

This is a nice reply if, like me, you like the historical-institutional account of philosophy. However, there are some important weaknesses. First, the argument fails to take into account what typically happens (and is supposed to happen) when a philosophy 'spawns' new disciplines. When, for example, psychology and economics moved away from philosophy and into separate disciplines, philosophers – after a while (and for the most part) – left them alone. They developed their own methods and training programs. Insofar as philosophers were interested in their results, they took those from the practitioners of those disciplines. Philosophers don't, qua philosophers, try to continue to contribute to the cutting edge of research in economics or psychology. In sum: the argument just outlined fails to take into account that there's now a separate discipline specializing in doing natural language semantics.⁵ When that's the case, historical continuity is broken and philosophers will eventually say: "We're done with those topics – they've been taken over by specialists."⁶

My prediction is that if this contingent historical connection is the *only* reason why philosophers still do natural language semantics, the work done by philosophers on natural language semantics is likely to slowly but surely come to an end. Of course, some people in some philosophy departments will continue to insist on hiring people to work on natural language semantics. They'll do this because they want people to talk to about these topics and don't realize (or wilfully ignore) that this work is better done in linguistics departments. So some philosophy departments will continue to hire philosophers of language who work on natural language semantics for a while, but in the long run this will stop, as it becomes clearer that semantics is best done elsewhere (in the same way it eventually became clear that economics and psychology were best done elsewhere).

3.3 Piecemeal and topic-specific significance of natural language semantics

I turn now to the third reply to the question "Why should those who reject the linguistic turn care about the properties of natural languages?" I call this Piecemeal and Topic Specific.

Piecemeal and Topic-Specific: the arguments for the linguistic turn were *too general* – their aim was to show something about the connection between natural language

⁵ In other words: We can't go straight from the observation that semantics is done in philosophy departments to the conclusion that semantics is philosophy because the observation doesn't tell the whole story. The whole story is that semantics is done in philosophy departments *and* in linguistics departments.

⁶ This is a version of The Division of Labor Argument above. Speaking loosely: To do cutting edge research on issues in e.g., economics, you need a PhD in economics. Without that, you're likely to be wasting your time. Of course, all of this is compatible with there being exceptions - either because of special skills of individuals or peculiarities of a topic.

semantics and philosophy *generally*. A better strategy is to look for piecemeal significances: not all of natural language semantics is philosophically significant, but here and there we find important connections and the challenge is to pin down these ‘here’s and ‘there’s. On this view, it’s too ambitious to expect a global argument against Abandonment. What’s more likely is to find a bunch of little points where there are interesting connections. So, for example, those working in epistemology have been interested in whether the verb ‘knows’ is context sensitive, and we can agree that answering that question might be important for responding to the skeptic and to other core issues in epistemology. Those working on identity statements have been interested in rigidity. Those interested in modality have been interested in features of modal language. These are all piecemeal interesting connections. So that’s a bit of evidence that in each of those cases there are philosophical questions that can be illuminated through the study of natural language semantics.

I suspect this is an underlying thought for many of those philosophers who practice natural language semantics qua philosophers, but it shouldn’t convince proponents of Abandonment. Here is why:

Piecemeal Not Good Enough: Piecemeal significance isn’t good enough because there’s piecemeal significance for *very many disciplines* and we don’t think natural language semantics should be a part of those other disciplines.

All kinds of issues about the interpretation of language come up in law, but that’s not a good argument for law faculties to hire people with law or jurisprudence degrees who work on natural language semantics. Their work should be informed by work in linguistics, but it wouldn’t make sense for those trained in law and jurisprudence to do the work of linguists. Similarly, the structure of natural languages is important in large swathes of history, psychology, literature, archaeology, and mathematics, among other disciplines. More generally, it’s hard to think of any discipline where facts about the semantics of natural languages don’t have *some* piecemeal significance. But if you agree that piecemeal significance in those domains provides no reason for those departments to hire people working on natural language semantics, why should piecemeal significance do that for philosophy? The flipside of this argument is that there are lots of issues that are of piecemeal significance for philosophy, but where we don’t even have the beginning of a good reason for studying them qua philosophers. E.g, Fermat’s Conjecture is important for some philosophical issues, but proving Fermat’s Conjecture is a job for mathematicians, not philosophers.

4. Williamson on the significance of argument structure and natural language semantics

I turn now to Williamson's defense of the claim that the study of linguistic structures is more important in philosophy than in other disciplines. Despite his attack on the linguistic turn in The Philosophy of Philosophy, he does not advocate the Abandonment strategy outlined above. Why not? The answer has to do with *the importance of clear argument structure in philosophy*.

Williamson says:

The paradigms of philosophical questions are those that seem best addressed by armchair considerations less formal than mathematical proofs. The validity of such informal arguments depends on the structure of the natural language sentences in which they are at least partly formulated, or on the structure of the underlying thoughts. That structure is often hard to discern. We cannot just follow our instincts in reasoning; they are too often wrong. [...] In order to reason accurately in informal terms, we must focus on our reasoning as presented in thought or language, to double-check it, and the results are often controversial. *Thus questions about the structure of thought and language become central to the debate, even when it is not primarily a debate about thought or language.* (2007: 45, my emphasis (HC))

The basic idea is clear enough: We philosophers rely extensively on informal arguments. In order to assess the validity of those arguments, we need to know the underlying structure of the (informal) sentences used to express those arguments. So it follows that 'questions about the structure' of those sentences become central in all parts of philosophy - in all those parts where we rely on informal arguments.

Reply to Williamson: The fundamental flaw in Williamson's argument is that it doesn't distinguish philosophy from any other discipline or indeed from anyone who is serious about thinking and arguing. Williamson is, of course, right that the validity of informal arguments depends on the structure of the sentences we use when reasoning.⁷ What is not true is the following assumption:

Underlying Assumption: If informal reasoning plays an important role in a debate, then questions about the structure of language are central to that debate.

This would overgeneralize massively. The argument could be applied to law, psychology, literature, politics, or any domain where participants argue in natural language and care about the quality of their arguments. That feature is in no way distinctive of philosophy. Indeed, it's as true for two people having a serious conversation as it is for an academic. In sum: *There's nothing here that provides evidence of a distinctive significance for natural language semantics in philosophy.* The problem isn't just that endorsing Underlying Assumption would

⁷ In this paper I'm interested in the question of whether the study of natural language semantics should be part of philosophy and whether it is of more relevance to philosophy than it is to a number of other disciplines. I'm not here focused on questions about the structure of thought. Williamson, in *The Philosophy of Philosophy*, moves back and forth between talking about the structure of language and the structure of thought. My focus here is exclusively on linguistic structures. My view is that the structure of thought is best studied empirically by psychologists, but I won't argue for that claim here.

overgenerate. The problem is that the assumption is false. Many things can have relevance to a debate without being central to the debate. Logic is for example relevant to the debate about whether the UK should leave the EU, but logic is not central to that debate. More generally, issues having to do with language and logical structure are relevant to any argument, but not part of the argument.⁸

Elaborating on the passage above, Williamson uses a suggestive, but ultimately unhelpful, analogy. He says:

Some contemporary metaphysicians appear to believe that they can safely ignore formal semantics and the philosophy of language because their interest is in a largely extra-mental reality. They resemble an astronomer who thinks he can safely ignore the physics of telescopes because his interest is in the extra-terrestrial universe. In delicate matters, his attitude makes him all the more likely to project features of his telescope confusedly onto the stars beyond. Similarly, the metaphysicians who most disdain language are the most likely to be its victims.

Analytic philosophy at its best uses logical rigor and semantic sophistication to achieve a sharpness of philosophical vision unobtainable by other means. To sacrifice those gains would be to choose blurred vision. (2007: 46)

But this analogy isn't helpful if the aim is to show that questions about language are more important to philosophy than to other disciplines (more central to philosophical debates than to, say, the debate about Brexit). The analogy shows that those using telescopes should be interested in the nature of telescopes. The extension of the point is that those using language should be interested in the nature of language. But everyone uses language and so the proper conclusion is that *everyone* should be interested in language. The analogy provides no support whatsoever for a distinctive role for the study of natural language semantics within philosophy.

5. The roles of philosophy of language in philosophy: foundational issues and conceptual engineering

I turn now to the happy-face part of this paper. The conclusion so far is that it's hard to justify the current practice of philosophy journals publishing papers on natural language semantics – to do so is like having *Mind* or *The Journal of Philosophy* publish papers on how to diagnose PTSD or the inflation rate in Argentina. That's an unwise use of resources: it's better to leave that work

⁸ Could Williamson reply that philosophers are *more* interested in and concerned with validity of inference than others who argue? I think not. First, that's not a true generic – it's patronizing towards non-philosophical thinkers and talkers. Second, even if there are others who don't care as much about validity of their arguments, then they *should* be more concerned with it. As a matter of fact, this is crucial to their work, even if those doing the work don't recognize it.

to those with proper training. However, none of this is to say that there's no role for philosophy of language in philosophy. I see two roles, one very obvious and another less obvious but more important.

5.1. First role for philosophy of language in philosophy: foundational issues in linguistics

There's a well-established practice of philosophers working on foundational issues in other disciplines. Philosophers work on the foundations of biology, physics, mathematics, and the social sciences. In general, for any discipline, D, there are philosophical questions that concern the foundations of D.

What are foundational questions about a discipline? I won't attempt a precise definition of 'foundational' here: for the purposes of this paper I'll treat the notion as a primitive and take it that readers have at least a partial grasp of that notion. Here are some rough diagnostics:

Rough diagnostics: the answers to foundational questions for a discipline will affect the interpretation of very many of the results in that discipline; they are questions that the methods of the discipline aren't particularly well suited to settle; they are questions about which there's typically no consensus among practitioners of that discipline.

These aren't necessary and sufficient conditions, but they are helpful generic claims about foundational questions. Here is a paradigm instance of a foundational question:

A paradigm: The question about whether numbers exist and if so how we can have knowledge about them is a foundational question in mathematics: the answer to that question will significantly affect the interpretation of many parts of mathematics (if, for example, Hartry Field (1980) is right, all of mathematics is false), it's not a question standard mathematical methods are well suited to answer, and there's no consensus among mathematicians about the answer.

What counts as 'foundational' in a particular discipline will vary over time and will be sensitive to the evolution of research methods and interests of the practitioners. So there will, at least to some extent, be constant evolution in what falls into the category of 'foundational'.

Turning now to the relationship between philosophy of language and linguistics, the remit of the philosophers will, as in other cases, include foundational issues in linguistics. What's included in that? Here are some candidates:

- What it is for something to have representational capacities?
- What are properties and situations (insofar as they play a role in semantics)?
- What are types (insofar as they play a role in, say, syntax or phonology)?
- What are languages, metaphysically speaking? For example: are languages abstract objects or psychological entities or something else altogether?

- What is truth (insofar as this affects e.g. semantics)?
- What is successful communication?
- What is agreement and disagreement?
- What is the relationship between language and thought?

Note that doing just what semanticists do – i.e. proposing semantics for particular fragments of a natural language – is not included in this list. And what is included in the list will vary over time, depending on how linguistics and the interests of its practitioners develop.

This description leaves a number of issues unsettled. Suppose that at a certain point in time not many linguists work on metaphor. Metaphor isn't exactly a foundational question about linguistics, but it's an important topic, there's a long philosophical tradition of working on it, and there's no clear methodology that linguists would agree on for investigating metaphor. What I've said above gives fairly little guidance in such cases and I think that's the right result: It would be silly to aim for a decision mechanism for how to distribute work between disciplines. There will be a broad range of undecided cases and, of course, also some that overlap.

5.2. Second role for philosophy of language: Conceptual engineering and normative reflections on language

The role of philosophy of language sketched in the previous section is unsurprising: it's an instantiation of a ubiquitous pattern of interaction between philosophy and other disciplines. The role I now turn to is specific to philosophy of language and I think more important and interesting.

Here is how I'll proceed:

- (i) First I'll describe conceptual engineering as I understand it.
- (ii) I'll then provide some history: I'll argue that this is what philosophy of language originally was, and focusing on these issues is a return to the origins not only of philosophy of language but also of analytic philosophy.
- (iii) At the end I'll explain why a renewed focus on conceptual engineering will, in effect, again make philosophy of language a form of first philosophy. Or, otherwise (and less provocatively) put, this focus will make philosophy of language central to very many other parts of philosophy.

5.3. Conceptual engineering: assessing and improving our representational devices

Conceptual engineering, as I understand it, has two parts:

- The first part involves an effort to develop a theory of how to assess and improve our representational devices.
- The second part involves practicing the theory developed in part one, i.e., assessing and improving particular representational devices.

If you think the primary representational devices should be called ‘concepts’, then this is (i) the effort to provide a theory of conceptual deficiencies and strategies for how concepts can be improved, and (ii) the effort to find deficiencies in specific concepts (such as truth, knowledge, freedom, identity, justice, gender and race) and then propose ameliorative strategies for those specific concepts. Following Blackburn (1999) and Scharp (2013), I call this kind of project ‘conceptual engineering’.

With this characterisation in mind, here are some salient questions for those working on conceptual engineering:

- What are the various ways in which representational devices can be defective?
- What are the various ways in which representational devices can be improved?
- What are the limits of revision, i.e., how much improvement is too much?
- What are practical strategies for improvement, i.e., how can an ameliorative project be implemented?
- How does a theory of conceptual engineering fit into an overall theory of language and mind?
- How much control do we have over conceptual change?⁹

Then in addition to these very general questions, there are important questions about how to conceptually engineer specific concepts. Consider the following cluster of questions: Are there deficiencies and ameliorative strategies for any of the following concepts: truth, freedom, race and gender concepts, moral concepts? For all of these, there are philosophers who answer the first conjunct of the question ‘yes’ and then go on to propose ameliorative strategies.¹⁰

5.4. A brief history of 20th century philosophy: conceptual engineering as the origin of philosophy of language and analytic philosophy

The view that philosophy of language should have as its primary focus conceptual engineering should not be seen as surprising or revisionist or in any way radical. This is what philosophy of language was always meant to be. It is the idea that the field should be consumed with efforts to describe the semantics of natural languages that should strike us as bizarre (and surprising and

⁹ For more on conceptual engineering, see e.g., Burgess and Plunkett (2013a, 2013b), Eklund (2014), Ludlow (2014). Cappelen (forthcoming) presents and expands on all the issues mentioned above.

¹⁰ See for example, and respectively: Scharp (2013), van Inwagen (2008), Appiah (1992,1996), Railton (1989), the papers collected in Haslanger (2012).

revisionist). In *Fixing Language: An Essay on the Foundations of Conceptual Engineering*, I sketch a brief history of 20th century philosophy to make a case for this claim. It's a version of this history that's very different from what's found, for example, in Scott Soames's two volumes on that topic (2003a, 2003b). As Soames sees it in that work, the normative (engineering) aspects played a marginal role in the origins of analytic philosophy and philosophy of language. The central aims of the analytic tradition were all descriptive (to describe various features of our linguistic practices and conceptual structures). That is, I think, a fundamental misunderstanding of philosophy in the 20th century and of the origins of analytic philosophy. The most crucial characteristic of philosophy in that period was the normative and engineering proposals: those were what initiated and motivated what we now describe as 'analytic' philosophy. To make that case would, of course, go beyond the scope of this paper, so I will just sketch some evidence for that way of looking at things.

First, consider Frege. One of the central aims of his work was to improve on what he took to be deficiencies of natural language. One can read all of Frege's *Begriffsschrift* as an exercise in conceptual engineering. Frege says:

If the task of philosophy is to break the domination of words over the human mind ..., then my concept notation, being developed for these purposes, can be a useful instrument for philosophers (Frege 1889: 7)

He clearly endorses the antecedent of the initial conditional: the task of philosophy (or as we would put it today: philosophy of language) is to break the domination of words over the human mind. The words of ordinary language have a negative effect on the human mind and Frege's aim was to help overcome this.

Second, significant elements of Wittgenstein's work were an effort to describe ways in which language could be defective (meaningless or 'on holiday') and then find ways to improve it.

Third, with Carnap (and the logical positivists) this focus on conceptual engineering became even more pronounced. Consider the following three aspects of Carnap's work that are related to engineering:

- First, the process he calls 'explication' is an effort to improve our conceptual apparatus. The idea is to take a term that suffers from various deficiencies (Carnap's focus was on what he calls 'vagueness' and 'indeterminacy') and then improve it along various dimensions.
- Second, according to Carnap, a great deal of speech is nonsensical because it doesn't satisfy the conditions for being meaningful. This is maybe the most fundamental kind of representational defect: the failure to be representational at all. In the paper 'The Overcoming of Metaphysics Through the Logical Analysis of Language' (1932), Carnap tries to show that most of Western philosophy and large swathes of ordinary speech are, literally, meaningless. Nonsensical speech is bad, he thinks, and should be avoided. So the goal of his paper is normative: point out ways to improve speech and thought.

- Third, and more controversially, we can understand Carnap's proposed theory of meaning – the verification theory – as a normative suggestion.¹¹ That is, it is a theory we should adopt, not because it is an accurate description of what meaning *is*, but because this would be a good thing for meaning to be. It would improve our speech and talk if this was what counted as meaningful. So, on this view, the verification theory is a theory of what meaning should be.

Finally, note that the so-called 'ordinary language' movement that dominated much of philosophy in the second half of the 20th century also had an important normative engineering-related component. In 'A Plea for Excuses', Austin says that '... ordinary language ... embodies ... the inherited experience and acumen of many generations of men. ... If a distinction works well for practical purposes in ordinary life (no mean feat, for even ordinary life is full of hard cases), then there is sure to be something in it, it will not mark nothing' (Austin 1956: 11). This is practical advice about the usefulness of various distinctions we make in ordinary language. Austin then notes that '...ordinary language is not the last word: in principle it can everywhere be supplemented and improved upon and superseded.' (Austin 1956: 11). The challenge here is to recognize when ordinary language is good enough and when it can be improved upon. This is a deeply normative project, not a primarily a descriptive one, and it is continuous with the kind of engineering projects described above.

Something weird and unfortunate happened around 1970: much analytic philosophy of language took a sharp descriptivist turn and the various engineering projects were downplayed. In philosophy of language, this was reflected in closer collaboration between philosophers of language and linguists. It might have been motivated, in part, by the perception that many of those who saw themselves engaged in various engineering projects did not have enough understanding of what they were criticizing. In order to criticize and improve something, you first need to understand how that thing works. So the descriptivist turn might, in part, have been motivated by a sense that the engineers needed more facts about language to work with. Whatever the reason,¹² that was the period when the descriptive/not-normative/not-engineering focused philosophy of language started to push the normatively focused philosophy of language into the periphery. (The focus on the groundbreaking work by Kripke, Putnam and Burge is maybe the clearest illustration of this change in direction.)

The descriptivist turn wasn't limited to those working in philosophy of language – it spread across philosophy. Epistemology is the clearest example. Since the publication of Gettier's paper in 1963, an almost unbelievable number of papers and books have been devoted to describing the exact way in which people who speak English use the word 'know', and to trying to understand

¹¹ I say we can so understand it, not that Carnap so understood it – though there's some evidence that Neurath understood it in this way. See Carnap 'Intellectual Autobiography', p23, in Schlipp (1963).

¹² Figuring out just what happened at this time and what motivated what I think of as 'the descriptivist turn in philosophy', will, I predict, be a central challenge for those trying to write the history of philosophy in the 20th century. If the rough outline I have sketched here is correct, then this was a pivotal, but poorly understood, juncture in the evolution of philosophical thought.

why some of them don't like to apply the term 'know' in so-called Gettier cases. This purely descriptive turn in epistemology is a historical aberration, as Robert Pasnau has documented (Pasnau 2013). Traditionally, what was called 'epistemology' was normative, concerned with describing epistemic ideals and strategies for improvement rather than engaging in careful lexicology.

The picture just sketched is incomplete. The normative tradition was kept alive but it was somewhat sidelined and marginalized. There are still many examples of prominent philosophers working on conceptual engineering: much contemporary work in feminist and race theory is conceptual engineering (see, for example, the papers collected in Haslanger 2012, Appiah 1992, 1996, Mallon 2004, 2006), a sub-literature on paradoxes is conceptual engineering (notably Scharp 2013), and in moral philosophy there are various conceptual engineering projects (e.g. Railton 1989 and the revolutionary fictionalism of Joyce 2005). For a more thorough overview, see chapter 2 of my *Fixing Language* and Burgess and Plunkett (2013a, 2013b). However, this kind of work is now seen as a kind of side-project, somewhat marginalized from the mainstream – it's not the core project. What I'm proposing is that we return to the roots of analytic philosophy and make this the core of philosophy of language and, I argue below, in some sense of philosophy more generally.

5.5. Final and central point: Philosophy of language construed as conceptual engineering is of central importance to all parts of philosophy

What would be the effect of refocusing philosophy of language on conceptual engineering, i.e. of returning philosophy of language to its roots? There are three important effects:

1. *It helps draw the distinction (the division of labor) between philosophy of language and linguistics*

The task of the linguist is to describe features (phonological, syntactic, semantic, morphological, pragmatic) of various natural language expressions. The task of the conceptual engineer is to assess our representational, including linguistic, devices and develop ameliorative strategies. This is one of the central tasks of the philosophy of language. The other is to work on foundational questions in linguistics in the way philosophers of physics work on foundational issues in physics (without doing physics) and philosophers of mathematics work on foundational issues in mathematics (without contributing to research in mathematics).

2. *Philosophy of language so construed is of great importance to all branches of philosophy*

Three assumptions have been central to the outline of conceptual engineering above:

(i) There can be a general theory of what conceptual engineering is and developing such a theory is or should be the central aim of philosophy of language.

(ii) Various concepts are central to different subdisciplines of philosophy: beauty, knowledge, justification, belief, perception, law, consciousness, etc. are all central to various part of philosophy.

(iii) For all these one can and should ask: *Is this concept defective and if so, can it be improved?*

If those three assumptions are true, then the general theory of conceptual engineering is of overarching significance for all of philosophy. Using some old-fashioned terminology, one could say that this makes philosophy of language – construed as conceptual engineering, not as natural language semantics – a form of first philosophy. Or at least this much is true: in order to do any other part of philosophy, you need to also do philosophy of language.

3. The relevance of philosophy of language to other intellectual disciplines, public discourse, and informal conversations

Analogous points can be made about other domains of speech and talk:

- The significance of philosophy of language – construed as an effort to develop a general theory of conceptual engineering – isn't restricted to philosophy. It's relevant to *all* intellectual disciplines: if the core concepts in mathematics, economics, biology, law or psychiatry can be defective and be improved, then practitioners of these disciplines should also be interested in – and informed by – a general theory of conceptual criticism and amelioration.
- Moreover, public discourse can and will often be conducted using defective concepts. Consider, for example, concepts such as person, hacker, rape, immigrant, refugee, marriage, and illness. These are examples of concepts that are central to important public debates. If those concepts are defective, then those debates are defective. A general theory of defective concepts can help us move those debates forward.
- Finally, when a single person is engaged in reflection about her life – about what she wants to do and what she wants to be – then she may be using defective concepts and amelioration will again be both relevant and important.

The point above is that philosophy of language – construed as in this paper – is of relevance *to anyone who thinks or talks and who wants to do so as well as possible*. If conceptual defects are few and far between, then this is not a very urgent relevance – it only alerts you to a distant possibility. If conceptual defects are ubiquitous, then the relevance is much more urgent: it should be at the top of everyone's to do list. I think defects are ubiquitous, but arguing for that will require a theory of conceptual engineering – *Fixing Language* tries to make that case.

6. Concluding remarks: two complications/reservations

I end with two complications/reservations.

Reservation/Complication 1: Don't Rock the Boat

In part this paper has been sociological: it's about how to divide labour between academic disciplines. Now you might think that top-down efforts to impose structures of that kind are misguided. Here is a natural thought:

Natural Thought: if good work on topic X is done by people in department D, why rock the boat? Who cares what department it is done in as long as it is done? If there's a good community of people finding answers to X-related questions, let them be and support them – don't impose structure from above. That kind of policing is bound to be counterproductive.

This is a thought I'm sympathetic to and the paper isn't a proposal to current university administrators about serving eviction notices to people who happen to work on semantics and have offices in philosophy departments. Think of the proposal in two ways.

First, it's a description of a rational division of labor between linguistics and philosophy – but of course actual institutions will be non-ideal (irrational) along many dimensions and also be sensitive to specific historical and contextual contingencies.¹³

Second, think of this also as a description of the connections between the cluster of questions we call 'philosophical'.¹⁴ It matters how we see the relevance of, and interconnections between, those questions. The linguistic turn that dominated much of 20th century philosophy saw questions about language as enormously significant. That view is now abandoned by most practitioners of philosophy and so it's an interesting question how to think of the connection between issues in semantics and the rest of philosophy. I've argued for a threefold conclusion: (i) questions about natural language semantics are no more central to philosophy than they are to any other intellectual discipline (the only exception being linguistics, where they are among the core questions), (ii) foundational questions about linguistics fit a pattern of philosophical work on foundations of various disciplines, (iii) philosophy of language construed as conceptual engineering is (always has been) of central significance to all parts of philosophy (and, more broadly, to anyone who is engaged in thinking and talking).

Complication/Reservation 2: This shows we need to revise institutional structures

¹³ One such contingency could be: these people are now doing good work, so leave them alone for awhile.

¹⁴ And that's a different question from how to think about the organisation of particular academic institutions.

Finally, I've been talking about how to divide work between what can roughly be described as 'academic disciplines' where these are individuated in fairly traditional ways. I have, for example, appealed in my arguments to how such divisions have evolved in the past (I appealed, for example, to what happened when psychology and economics separated from philosophy). There's a kind of conservatism built into that way of thinking. It assumes that we need these divisions and that academic work should in some significant way be guided by them. You might think the right reaction to the kind of considerations above is a form of institutional revisionism:

Institutional Revisionism: The kind of work semantics exemplifies shows that academic institutions must become less rigid and rely less on traditional divisions like 'philosophy' and 'linguistics'. When these distinctions are loosened up, the question: "Does it belong to philosophy or to linguistics?" become the wrong kind of question to ask, because those are the wrong kind of categories to use when thinking about intellectual work.

Again, I have sympathy with this way of thinking. It might be the right reaction to the arguments in this paper, but I have to leave an exploration of that option to others. It's important and intellectually challenging to explore revisionist ideas about academic structures, but one needs *some* kind of structure and coming up with a concrete alternative is a massive task – one that no doubt will be met with equally massive institutional and administrative resistance. The actual implementation of any revision is very, very far away. In the meantime we need to figure out how best to work within the structures we have.¹⁵

Bibliography

Appiah, Kwame Anthony (1992). *In My Father's House: Africa in the Philosophy of Culture*. Oxford University Press.

Appiah, Kwame. Anthony (1996). "Race, culture, identity: misunderstood connections", in Peterson, Grethe B., *The Tanner lectures on human values XVII*, Salt Lake City: University of Utah Press, pp. 51–136

Austin, John (1956). A plea for excuses, *Proceedings of the Aristotelian Society* 57:1--30.

Ayer, A. J. (1936). *Language, Truth and Logic*. Victor Gollancz, Ltd.

Blackburn, Simon (1999). *Think: A Compelling Introduction to Philosophy*. Oxford University Press.

Cappelen, Herman (2013). *Philosophy Without Intuitions*. Oxford University Press

Cappelen, Herman (forthcoming). *Fixing Language: An Essay on Conceptual Engineering*

¹⁵ Thanks to Alexi Burgess, Robyn Carston, Josh Habgood-Coote, Matthew McKeever, Dilip Ninan, and Seth Yalcin for helpful comments.

Carnap, Rudolph (1932). Überwindung der Metaphysik durch logische Analyse der Sprache, *Erkenntnis*, 2: 219–241

Dummett, Michael A. E. (1978). *Truth and Other Enigmas*. Harvard University Press.

Eklund, Matti (2014). Replacing Truth, in *Metasemantics: new essays on the foundations of meaning*, Alexis Burgess, and Brett Sherman (eds), Oxford: OUP: 293-310

Field, Hartry (1980). *Science Without Numbers*. Princeton University Press.

Haslanger, Sally (2012). *Resisting Reality: Social Construction and Social Critique*. Oxford University Press

Frege, Gottlob (1889). *Begriffsschrift, eine der arithmetischen nachgebildete Formelsprache des reinen Denkens*, Halle a. S.: Louis Nebert. Translated as *Concept Script, a formal language of pure thought modelled upon that of arithmetic*, by S. Bauer-Mengelberg in J. van Heijenoort (ed.), *From Frege to Gödel: A Source Book in Mathematical Logic, 1879–1931*, Cambridge, MA: Harvard University Press, 1967.

Joyce, Richard (2005). *The Myth of Morality*. Cambridge University Press.

Kant, I. *Groundwork of the Metaphysics of Morals*, Mary Gregor and Jens Timmerman (trans) CUP: Cambridge.

Ludlow, Peter (2014). *Living Words: Meaning Underdetermination and the Dynamic Lexicon*. Oxford University Press.

Mallon, Ron (2004). Passing, traveling and reality: Social constructionism and the metaphysics of race. *Noûs* 38 (4):644–673.

Mallon, Ron (2006). 'Race': Normative, not metaphysical or semantic. *Ethics* 116 (3):525-551

Pasnau, R. (2013). Epistemology Idealized. *Mind* 122 (488):987-1021.

Railton, Peter (1989). Naturalism and Prescriptivity. *Social Philosophy and Policy* 7 (1):151.

Rorty, Richard (1967). *The Linguistic Turn*. Chicago, University of Chicago Press.

Scharp, Kevin (2013). *Replacing Truth*. OUP

Schlipp, Arthur, ed (1963). *The Philosophy of Rudolf Carnap*. Open Court Publishing.

Soames, Scott (2003a). *Philosophical Analysis in the Twentieth Century Vol 1: The Dawn of Analysis*. Princeton University Press.

Soames, Scott (2003b). *Philosophical Analysis in the Twentieth Century Vol. 2: The Age of Meaning*. Princeton University Press.

van Inwagen, Peter (2008). How to Think about the Problem of Free Will. *Journal of Ethics* 12 (3/4):327 - 341.

Williamson, Timothy (2007). *The Philosophy of Philosophy*. Blackwell.