

Artificial Language Philosophy of Science

Sebastian Lutz*

Draft: 2011-03-13

Abstract

Artificial language philosophy (also called ‘ideal language philosophy’) is the position that philosophical problems are best solved or dissolved through a reform of language. Its underlying methodology, the development of languages for specific purposes, leads to a conventionalist view of language in general and of concepts in particular. I argue that many philosophical practices can be reinterpreted as applications of artificial language philosophy. And many factually occurring interrelations between the sciences and philosophy of science are justified and clarified by the assumption that an artificial language methodology is applied in both.

1 Introduction

As there are different methodologies in philosophy, so are there different methodologies in the philosophy of science. This article aims to articulate and defend one specific methodology, which has been alternately criticized and ignored in recent times: artificial language philosophy. According to this view, philosophical problems are best solved or dissolved by the development of new languages and concepts and the regimentation of existing concepts (Rorty 1967a, §2), also called their “improvement” (Carnap 1963, §19), “reform” (Maxwell and Feigl 1961), or “explication” (Carnap 1950b, §§2–5).¹

*Theoretical Philosophy Unit, Utrecht University. sebastian.lutz@gmx.net. An earlier version of this article was presented at the *Future of Philosophy of Science* conference at Tilburg University, the Netherlands, in 2010. I thank the participants, Thomas Müller, and two anonymous referees for very helpful comments. I thank Alana Yu for excellent suggestions for improving the clarity and style of the article.

¹Rorty (1967a) popularized the term ‘ideal language philosophy’, but as Matteo Collodel, Eric Schliesser, and an anonymous referee have pointed out, the name suggests that there is a unique ideal language. The term’s originator, Bergmann (1949, 439), similarly assumed that there can be one ideal language for all contexts. Since mainly expositors (e. g. Rorty 1967a; Lutz 2009) have applied the term to the works of Carnap and philosophers of a similar inclination, while Carnap (1963, 938) himself speaks of the construction of “artificial languages”, the term ‘artificial language philosophy’ seems more apt.

Neither ignored nor as widely criticized is the methodology known as ordinary language philosophy, which is based on the idea that philosophical problems are best solved or dissolved by investigating ordinary language (Rorty 1967a, §2; Kauppinen 2007). Even more widely accepted as a philosophical methodology but not as easily circumscribed is naturalized philosophy, roughly the position that philosophical problems are best solved or dissolved through empirical research (Giere 1985; Feldman 2008, §2). Probably even more widely accepted, and even less easy to circumscribe, is what I will call, for lack of a better term, ‘traditional philosophy’. In the paradigmatic case, traditional philosophy shares with empirical research the aim of arriving at truths about the world, but without direct recourse to empirical methods (cf. Cohnitz and Häggqvist 2009, 9).

These four philosophical methodologies, though probably neither exhaustive nor mutually exclusive, clearly diverge in important ways. The problem of the relation between causation and explanation, for example, would be addressed in artificial language philosophy by suggesting concepts of causation and explanation, while the approach in ordinary language philosophy would be to analyze how the two terms are used in ordinary language. A naturalized approach to the problem would be an empirical investigation into the development of the concepts. A traditional approach might, for instance, use intuition as a guide for better understanding the nature of causation and explanation. It would be surprising if all of these methodologies led to the same results, where a result is a solution or dissolution of a problem in a way that is satisfactory within the respective methodology.

It would also be surprising if all of these methodologies led to a philosophy of science whose results are fruitfully connected to the sciences. While it is open to debate whether ordinary language and traditional philosophy fulfill this desideratum for a philosophy of science, naturalized approaches clearly do. But it is further a desideratum for any philosophical methodology that it address philosophical problems, and naturalized approaches to philosophy have often been charged with simply changing the topic. In this article, I will argue that artificial language philosophy of science does fulfill both desiderata, and I think that this amounts to a pragmatic defense of this methodology.

There are different ways to argue against a methodology. It may have too little content to ever yield a result (a charge that, for example, William Craig once leveled against Feyerabend’s credo “Create more theories!” (Demopoulos 2008, §9)). Or a methodology might have too much content and yield mutually inconsistent results. It might also lead to false empirical claims. A conclusive defense of any methodology, it seems, would need to establish that the methodology always leads to true results. Such an argument is hard to come by in philosophy because typically, philosophical claims are neither empirically testable nor straightforward theorems of some formalism.

There are many discussions about philosophical methodology (e. g., Williamson 2007; Papineau 2009), some of which even take into account artificial language philosophy (Rorty 1967b; Lutz 2009). But, arguably, no specific methodology has been conclusively

shown to be correct, and of the ones discussed in this article, none has been conclusively shown to be wrong. In such a case, a cumulative strategy is often pursued: The individual advantages of each methodology are compared and, taken together, provide a defense of one of them. For such a cumulative strategy, pragmatic arguments become relevant (Feigl 1969). One such pragmatic argument contends that a certain philosophical methodology leads to results in philosophy that are neither inconsistent nor empirically false (Rorty 1967a, 3). My point, on the other hand, is that artificial language philosophy leads to results in meta-philosophy (that are neither inconsistent nor empirically false).

It is important to note that the role of this article's meta-philosophical results in such a cumulative defense of artificial language philosophy should not distract from the content of the results themselves. First and foremost, the article lays out one philosophical methodology's meta-philosophical implications for the status of philosophical results and the relation of science and philosophy of science.

The rest of this article is structured as follows: After a short overview of the relations between traditional, ordinary language, and naturalized philosophy (§2), I provide an outline of artificial language philosophy and its relation to the sciences (§3). I then argue that artificial language philosophy can capture much of philosophical practice that ostensibly follows traditional, ordinary language, or naturalized philosophy (§4), which suggests that artificial language philosophy addresses philosophical problems. Finally, I show how many relations that have been observed between science and philosophy of science can be justified and clarified within artificial language philosophy (§5).

2 Relations between philosophical methods

2.1 Traditional philosophy

Traditional philosophy can be considered an investigation of facts about the world. However, a straightforward empirical investigation of, say, hydrochloric acid differs from the philosophical investigation of causation in that hydrochloric acid has a specific density, decomposes at a specific temperature, and generally has properties that can be determined with a certain degree of confidence through measurements and, ultimately, through observations of the outcomes of these measurements. And there are experimental ways to determine whether an unknown substance is hydrochloric acid. On the other hand, there is no measurement to determine how causation connects to explanation, and there is no experimental way of determining whether a specific situation or process is an instance of explanation or causation.

In traditional philosophy, the analogue of experiments and observations are often intuitions (cf. Sosa 2007, 105). Papineau (2009, 14), for example, states that

Gettier appealed to the intuition that a belief whose truth is accidental rel-

ative to its method of justification is not knowledge; Kripke appealed to the intuition that something that is not the causal origin of a name is not its bearer; and so on. On my account, all these intuitions are synthetic claims about the relevant kind of scenario.

In other words, these intuitions are about the world. Feigl (1958, 6) distinguishes between two kinds of intuitions: “hunches”, which can be tested by observations, and “trans-empirical” intuitions, which cannot. If traditional philosophy were to rely on hunches, it would be amenable to empirical test. However, according to Papineau (2009, 18) it relies on trans-empirical intuitions:

If my judgemental procedures decide who is a knower by assuming, *inter alia*, that accidentally true believers are not knowers, then clearly there isn't any question of my meeting up with a case where I judge such an accidentally true believer to be a knower after all. Again, if my judgmental procedures decide what thing bears some name by noting the causal origin of the use of the name, then I'm not going to come across cases where I judge that some name is borne by something other than its causal origin. But this impossibility of direct falsification does not mean that the relevant general assumptions are analytic. They may yet have a substantial synthetic content [. . .].

Both for Papineau (2009, §IV) and Williamson (2007, §6), thought experiments are a core method in philosophy. But while Papineau considers thought experiments to elicit intuitive judgments, Williamson considers them to elicit judgments simpliciter. This is because Williamson (2007, 3) holds that “so-called intuitions are simply judgments (or dispositions to judgment)”, which, however, brushes over an important distinction: Some judgments can be supported by explicit argument or observation, while others cannot be, or at least not completely. These judgments that do not have sufficient support are intuitive, and they are used analogously to observations. Williamson (2007, §7.2) gives a short overview of other philosophers' positions on intuition; in all of them, intuitions are used as premises, not as conclusions.

It is clear that intuitions about, say, the presence or absence of causation or explanation in individual thought experiments cannot entail a general rule about the relation of causation and explanation. Such a rule either has to be postulated like a scientific hypothesis, or gathered from more general intuitions about the relation itself. The general rule can be tested against the intuitions in individual thought experiments, and in case of inconsistency, the intuitions about the general rule or about the thought experiments have to be modified. This method of testing and revision may be repeated and may eventually lead to a reflective equilibrium, where the intuitions about the general rule and the thought experiments agree (Daniels 2008, §1). The test for consistency is, in a sense, a purely formal method that is used in every philosophical methodology here discussed. How an inconsistency is resolved, however, will rely crucially on the intuitions themselves.

Feigl (1958, 7f) also distinguishes between an intuition and its target, where the intuition is a psychological phenomenon and its target is some fact about the world. Strictly speaking, it is then not the intuition that features in a philosophical argument, but a description of its target or, following the terminology by Fedyk (2009, §2), the propositional content of the intuition. When intuitions are considered as psychological phenomena, it is of interest how their occurrence can be established (Feigl 1958, 8–11). If the intuitions under examination belong to a specific class of people (e. g., some group of philosophers, ordinary people, scientists), then the best method of determining the content of those intuitions seems to be a statistical one. The relevant intuitions may also be those that a specific class of people *would* have, if presented with some class of facts, considerations, and examples.² Then the best and probably only method of determining the intuitions' content is empirical psychology, which could establish that people in fact usually develop these intuitions. Some parts of traditional philosophy therefore may have to be naturalized.

2.2 Ordinary language philosophy

In traditional philosophy, intuitions are taken to provide information about the world; in the terminology of Fedyk (2009, §4), they are interpreted as world-directed. In ordinary language philosophy, on the other hand, the same intuitions are interpreted as providing information about the language in which the world is described; they are interpreted as meaning-directed. According to Williamson (2007, §1), this linguistic turn towards language has been largely superseded by a conceptual turn towards the concepts of thought. While I will discuss only ordinary language philosophy here, the discussion is also applicable to methodologies based on the analysis of concepts of thought. I will rather use 'concept' to refer to the intension of a term, such that a definition or some weaker set of meaning postulates for a term (e. g., 'cause') would determine the concept itself (e. g., *cause*).

If ordinary language philosophy relies on the actual language use of some group, or the actual linguistic intuitions of some group, then one difference between traditional and ordinary language philosophy is straightforward: The claims of traditional philosophy are about unobservable states of the world, while the claims of ordinary language philosophy are about observable states or, in the case of people's language intuitions, states which can be empirically determined with some certainty. Intuitions about people's linguistic intuitions and language use are therefore hunches in Feigl's sense. Of course, whether these hunches are accurate is itself an empirical question. Eventually, the truth of their propositional content has to be established empirically (Feigl 1958, 6f; Mates 1958, 165; Nadelhoffer and Nahmias 2007, 129; Sytsma 2010, §1). As in the case of trans-empirical

²This may be what Williamson (2007, 191, 216) has in mind when he speaks of philosophical judgments that require "philosophical training" leading to specific "skills".

intuitions, this also holds for intuitions about the language use or linguistic intuitions that people *would* have if presented with some class of facts, considerations, and examples.³ Some parts of ordinary language philosophy therefore have to be naturalized (Mates 1958).⁴

Like the use of thought experiments in traditional philosophy, the exclusive reliance on actual language use, that is, individual speech acts, cannot lead to the establishment of general rules about language use. Such a general rule may be either stipulated like any other empirical theory or gathered from more general intuitions about the rules of language. Once spelled out explicitly, the general rules may be inconsistent with the content of intuitions about individual speech acts. As in the case of traditional philosophy, such an inconsistency can be established through a purely formal analysis. If the method of reflective equilibrium is used to resolve such inconsistencies by excluding certain language uses or general rules, then ordinary language philosophy has a normative part (Carnap 1939, §4). In this case, any intuitions about how language *should* be used are trans-empirical, and the discussion in §2.1 applies.

3 An outline of artificial language philosophy

Like ordinary language philosophy, artificial language philosophy considers philosophical problems to be problems of language. Unlike ordinary language philosophy, however, artificial language philosophy contends that philosophical problems are best solved or dissolved not by the analysis of actual language use, but by the conventional prescription of a new language. The new language must be clear enough to make the original question amenable to formal proof (thus leading to a solution), or be such that the problem cannot be formulated (amounting to its dissolution) (Bergmann 1957, 326). An artificial language philosopher may also suggest a reform of language even for contexts in which there are no problems, either because the resulting language has some pragmatic advantage like greater simplicity or precision, or because this change helps to avoid problems in other contexts (Maxwell and Feigl 1961, 491).

In principle, there is no restriction on the choice of language. As Carnap (1934a, §17) states in his “Principle of Tolerance”, even the logic of a language is conventional. For what follows, however, it will suffice to outline a very simple case of language choice. I will assume that in addressing the philosophical problems at hand, the logic of the language is taken as given, and there is some set \mathcal{B} of *basic terms*⁵ whose application is taken as unproblematic. It may be taken as unproblematic because the terms apply more

³This may be what Kauppinen (2007, §5) has in mind when he claims that “(philosophical) dialogue and reflection” lead to a convergence of linguistic intuitions.

⁴Arguably, the naturalization leads to experimental philosophy (Lutz 2009, §3; Sytsma 2010, §§1–3).

⁵Here and in the following, a term is any non-logical constant of a language.

or less immediately to observations (cf. Chang 2005), but more generally, the basic terms simply refer to concepts that are not themselves under investigation when addressing the problems at hand (Reichenbach 1951, 49; Lewis 1970, 428). Since the application of basic terms is assumed to be unproblematic, \mathcal{B} -sentences (sentences containing only basic terms) can be established as true through empirical investigation. The choice of a language then amounts to concept formation, that is, the choice of meaning postulates for the terms not in \mathcal{B} . Let me call those terms ‘auxiliary terms’ and their set ‘ \mathcal{A} ’.

\mathcal{A} can contain terms for pre-theoretically understood concepts whose modification is intended to solve or dissolve the problems at hand. Such a modification is called ‘explication’, the pre-theoretic concept is called an ‘explicandum’, and the replacing concept the ‘explicatum’ (Carnap 1950b, §§2–5). The explicatum has to fulfill conditions of adequacy, which identify what problems the newly formed concept should solve or dissolve, and in what contexts it should be applicable. The latter are suggested by the pre-theoretic use of the explicandum (Kuipers 2007, §2). One may, to take a prominent example, consider the use of *causation*, and search for a new concept that applies only to pairs of events where the description of the first event provides an explanation of the second, and, if the first event is an action, the actor is responsible for the second event. If such a concept was found, it could be used instead of *causation*, and would relate in a clear way to *explanation* and *responsibility*. In this example, ‘event’, ‘explanation’, ‘action’, ‘responsibility’, etc. are assumed to be in \mathcal{B} , but each of these terms can be in \mathcal{A} in other contexts. In general, since the bipartition of terms into \mathcal{B} and \mathcal{A} is context dependent, a term P explicated in one context with the help of a term Q could in another context be used to explicate Q .

\mathcal{A} can also contain terms for entirely new concepts, which do not act as explicata. One may, for example, introduce concepts like *natural kind* or *possible world* simply because they solve some problem in philosophy, not because there are corresponding explicanda. In this case, the conditions of adequacy can be chosen freely and may, for instance, amount to a list of problems that the new concept should solve.

Because of the dependency on the problems a concept is meant to solve, the choice of the conditions of adequacy is ultimately pragmatic. The conditions of adequacy should be precise enough, however, to determine reliably whether they are fulfilled by a concept. Of two concepts that fulfill all of the conditions of adequacy, one may choose one over the other based on convenience, but in principle, they are equally adequate.

Carnap (1939, §24) gives a general outline for the semantics for a language bipartitioned into basic and auxiliary terms. Przełęcki (1969, chs. 5f) gives what can be seen as an elaboration of Carnap’s account in formal semantics, as do Suppe (1971) and Andreas (2010). I will briefly summarize Przełęcki’s account, since it essentially anticipates the other two, but is more developed and general. Because the basic terms are unproblematic, Przełęcki can assume that their interpretation is determined by a set \mathbf{B} of possible \mathcal{B} -structures. The interpretation of the auxiliary terms is then determined solely by \mathbf{B}

and by the sentences Π , the meaning postulates for the auxiliary terms, which have been chosen to be true in the process of concept formation. Thus the set \mathbf{M} of structures for *all* terms ($\mathcal{B} \cup \mathcal{A}$) contains all and only those models of the meaning postulates that expand the structures in \mathbf{B} to $\mathcal{B} \cup \mathcal{A}$; that is, $\mathbf{M} = \{\mathfrak{M} \mid \mathfrak{M}|_{\mathcal{B}} \in \mathbf{B} \text{ and } \mathfrak{M} \models \Pi\}$.

For ease of discussion, assume that the set \mathbf{B} is determined solely by empirical research, and not by meaning postulates for the basic terms.⁶ Then all \mathcal{B} -sentences are either logically true (that is, true in virtue of the chosen logic), or empirical (that is, determined by what the world is like). There are now three ways to determine the truth or falsity of a sentence φ involving auxiliary terms. All of them rely on formal proof, thereby sharing a core method with traditional and ordinary language philosophy. For one, it may be possible to prove the truth or falsity of φ from the meaning postulates for the auxiliary terms that occur in φ . If this is impossible, it may still be possible to prove that φ is true (or false) in all structures \mathbf{M} . Such a proof would rely on empirical research. Finally, if neither way is possible, new meaning postulates for the terms in φ must be developed so that one of the first two ways *becomes* possible. This last case thus involves concept formation, rendering the truth or falsity of φ a matter of convention. Since formal proofs rely on the meaning postulates and the logic, which have to be chosen by convention, there are only two methods of determining the truth or falsity of a sentence: language conventions (leading to Π) and empirical research (leading to \mathbf{B}). Since language conventions, unlike empirical results, are not about the world, analyticity thus amounts to conventionality. If philosophy does not engage in empirical research, this means that philosophical results are analytic, either because they are formal proofs within the language, or suggestions for language conventions. Note that formal proofs can be important for language choice, by making otherwise hidden features of a language obvious.

The difference between the first two methods and the last method of determining the status of φ illustrates the distinction between “internal” and “external” questions introduced by Carnap (1950a, §2). The internal questions are those that rely on a chosen language (in Carnap’s terms, a “linguistic framework”), that is, a chosen logic and a chosen set of meaning postulates. Within this language, the investigation of the status of φ is objective, and corresponds to the first and second methods: For, whether φ is true or false depends solely on the state of the world, \mathbf{B} , and the language, which is fixed.⁷ The third method, that of concept formation, is a means of answering external questions. Here, the truth or falsity of φ is not determined objectively, but rather by conventional choice, and a claim about the status of φ cannot be right or wrong, but only more or less practical. There is, in this sense, no fact of the matter.

⁶Carnap (1952) describes how to treat meaning postulates for basic terms, Przełęcki (1969, §10.II) gives a method for introducing meaning postulates, and Kyburg (1990) discusses a method of deciding between different sets of meaning postulates for basic terms in probabilistic theories.

⁷*Within* a given language, the main tenets of realism would therefore seem to be true (I thank an anonymous referee for this point).

The preceding considerations are also intended to apply to the languages of the sciences (Carnap 1966, 187f): Within a science, it is only possible to test empirically whether some auxiliary term applies in some situation if this can be inferred by formal proof based on **B** and Π . In all other cases, science has to engage in concept formation in the same way as philosophy.

Since a scientific theory is developed to accommodate empirical results, its language will have already been chosen for sake of describing and investigating these results. In fact, the notion of explication is explicitly modelled on concept formation in the sciences (Carnap 1950a, §3; Hempel 1952, §3). Accordingly, science is teeming with explicata, such as ‘temperature’ explicating ‘warm’ (Carnap 1950a, §4; Hempel 1952, §10), and completely new concepts like ‘phlogiston’, ‘oxygen’, and ‘gene’, introduced to account for phenomena that are described in basic terms. Therefore it is of interest to establish the rules of the scientific language precisely. Here the methodology of ordinary language philosophy can be very helpful, except that it has to be applied to scientific rather than ordinary language. To avoid exclusive reliance on linguistic hunches, the scientists’ actual usage of scientific concepts can be determined empirically (cf. Stotz et al. 2004). In this way, the application of the methodology of ordinary language philosophy would fulfill what Reichenbach (1938, 3) calls the “descriptive task of epistemology”, that is, it is a search for the rules of the scientific language that capture the language intuitions of the scientists. As Waters (2004, §3) argues however, even this descriptive task goes beyond pure observation, for actual usage is often too vague or inconsistent for the establishment of proper rules.

Reichenbach (1938, §1) identifies two additional tasks of epistemology. One is the “critical task”, the identification and evaluation of inferences. In the terminology of artificial language philosophy, this amounts to formal proof. Next is the “advisory task”, the proposal of concepts for use in the sciences. As Waters (2004, §§5f) lays out, Reichenbach sees the scientists as the final arbiters of language choice, but such a restriction is not inherent in artificial language philosophy in general. The conditions of adequacy are a pragmatic choice for the philosopher of science, and not every philosopher has the same goals as the scientists on whose research they rely.

4 Philosophical practice in artificial language philosophy

By design, artificial language philosophy is closely connected to scientific methodology. But it is exceedingly simple to develop a methodology with this feature. Naturalized epistemology (Quine 1969) and experimental descriptivism (Nadelhoffer and Nahmias 2007, §2) apply neuropsychology to the development and application of concepts, understood as neuropsychological phenomena. Naturalized philosophy of science (Giere 1985) applies sociology to theory choice, understood as a social phenomenon. It is straightforward

to see that such methodologies are closely connected to the sciences. It may be, however, that these methodologies simply change the subject—that they do not, in fact, address philosophical problems at all (cf. Kim 1988). The idea that naturalized methodologies cannot by themselves solve philosophical problems is also suggested by their possible use as a proper part of traditional and ordinary language philosophy.

A straightforward defense of artificial language philosophy against this charge would require a precise definition of ‘philosophical problem’ and ‘solution of a philosophical problem’. Since so far, neither term has been defined to general satisfaction, I only aim to show that artificial language philosophy succeeds in capturing many philosophers’ posited solutions to philosophical problems, that is, it captures much of philosophical practice. It is of course difficult if not impossible to investigate the entirety of philosophical practice, but I will argue that artificial language philosophy can at least capture many of the applications of the different philosophical methodologies described in this article.

Artificial language philosophy trivially captures its own applications. And in major discussions of the relation between artificial language and ordinary language philosophy (Carnap 1955, §1; Hare 1960, 158; Carnap 1963, §19), there is a general consensus that the results of ordinary language philosophy can be a starting point for the choice of a language (cf. Lutz 2009, 127f). However, an exclusive reliance on the language as-is would make for an inordinately weak philosophical methodology, for the problem of vague or inconsistent uses of ordinary (and scientific) language cannot be resolved by recourse to the very problematic uses in question. When increasing the precision of a language or resolving its inconsistencies, Maxwell and Feigl (1961, 494) argue, an ordinary language philosopher “*is actually indulging in tacit reformation and issuing a stipulation as to what the terms in question are to mean*”. Because of this normative part (which cannot obviously be naturalized), ordinary language philosophy therefore threatens to collapse into artificial language philosophy.

The practices of traditional philosophy can also be captured by interpreting alleged discoveries of facts as inventions of new concepts or whole new languages. In the terminology of Carnap (1934b, 13–17, 19), this means switching from the “material” or “connotative mode of expression” to the “formal mode of expression”. When intuitions are used as evidence, they are interpreted as meaning-directed rather than world-directed, and their successive development in the method of reflective equilibrium is interpreted as a method of explication (cf. Kuipers 2007, xiv). However, while traditional philosophy faces the challenge of justifying its claims as *discoveries*, and thus of explaining how philosophers gain cognitive access to those facts that the claims are about (e. g., the targets of their intuitions), artificial language philosophy can simply justify them as a prudent choices of language that avoid some problems and solves others.

Indeed, the descriptions of the methods of traditional philosophy by its practitioners sometimes already read like descriptions of artificial language philosophy. Reviewing a

critique of traditional metaphysics by Ladyman and Ross (2007), Dorr (2010) describes a tentative consensus among metaphysicians about methodology:

It is not enough simply to announce that Xs are more fundamental than Ys: if I want to defend this claim, I am supposed, at a minimum, to (i) introduce a language in which I can talk about Xs without even seeming to talk about Ys; and (ii) make some kind of adequacy claim about this language, e. g., that it can express all the genuine facts that we can express using Y-talk, or that all the Y-facts supervene on the facts stateable in the language. For example, if I want to maintain that spacetime is less fundamental than the spatiotemporal relations between bodies, I must describe a language for characterizing these relations, and explain how it can adequately capture, e. g., claims about the global topological structure of spacetime.

Furthermore, Dorr states that one “earn[s] the right” to consider a philosophical problem “dissolved [...] by describing a fundamental language within which no corresponding questions can be formulated”. If the “genuine facts that we can express using Y-talk” are taken to determine in which contexts the X-language should be applicable, Dorr essentially describes the conditions of adequacy on an artificial language as discussed in §3.

Dorr (2010) also emphasizes the importance of language choice compared to formal proof:

The whole approach [by Ladyman and Ross (2007)] reflects an exaggerated sense of the importance of argument in metaphysics, and a corresponding underestimation of the difficulty of merely crafting a view coherent and explicit enough for arguments to get any grip.

From the perspective of artificial language philosophy, this crafting of a “coherent and explicit” view is nothing but the search for a language in which philosophical problems can be solved by formal proof.

The practices of naturalized philosophy are hard to circumscribe because naturalized philosophy itself is hard to circumscribe, but to the extent that it complements traditional and ordinary language philosophy, artificial language philosophy can capture its practices as well. And to the extent that naturalized philosophy *relies* on empirical results rather than establishes them, it is engaged in language choice and formal work. When it thereby addresses philosophical problems, naturalized philosophy amounts to artificial language philosophy. The one kind of naturalized philosophy that artificial language philosophy cannot accommodate is empirical research into a non-linguistic phenomenon. A philosopher who empirically investigates the use of a term contributes to the description of the explicandum. However, a philosopher who, for instance, determines the angles of a triangle of light rays over great distances does not describe an explicandum or engage in

language choice or formal work. Such research, however, is often charged with not being philosophy at all.

As an illustration of the reinterpretation of philosophical practice in artificial language philosophy, consider how Sosa (2007, 104) responds to the claim by Nichols and Knobe (2007) that usage of ‘responsible’ in ordinary language is inconsistent because of a performance error:

In any case, there is an alternative explanation that will cast no affect-involving doubt on the intuitions in play. This other possibility came to mind on reading their paper, and was soon confirmed in the article on moral responsibility in the Stanford Encyclopedia of Philosophy, where we are told that at least two different senses of ‘moral responsibility’ have emerged: the attributability sense, and the accountability sense.[. . .]

So, here again, quite possibly the striking divergence reported above is explicable mainly if not entirely through verbal divergence.

For Sosa (2007, 100), the “use of intuitions in philosophy should not be tied exclusively to conceptual analysis. [. . .] Some such questions concern an ethical or epistemic subject matter, and not just our corresponding concepts”. This is presumably how he interprets the “emergence” of two kinds of responsibility: They both exist, but people refer to them with the same word, leading to “verbal divergence”.

Sosa probably refers to the fall 2004 edition of the *Stanford Encyclopedia of Philosophy*, in which Eshleman (2004, §2.2) writes that

at least some disagreements about the most plausible overall theory of responsibility might be based on a failure to distinguish between different aspects of the concept of responsibility, or perhaps several distinguishable but related concepts of responsibility.

Broadly speaking, a distinction has been drawn between responsibility understood as attributability and responsibility as accountability.

Eshleman’s formulation differs from Sosa’s paraphrase in that the disagreements stem from confusing two “different aspects of the *concept* of responsibility”. At least read literally, according to Eshleman it is the pre-existing concept (not a fact) that gets confused. An analysis of the concepts of ordinary language would resolve the inconsistency. Regarding such a line of reasoning, Maxwell and Feigl (1961, 489) note that there is little reason to think that the two aspects allegedly being confused are somehow already present in ordinary thinking. Certainly, the ordinary user of the concept is not aware of them—otherwise there would be no confusion. Even if the ordinary user were to agree with the clarification (the distinction between attributability and accountability), Maxwell and Feigl note, this agreement may very well be the result of a change of language.

According to artificial language philosophy, then, the introduction of the distinction between responsibility as attributability and responsibility as accountability into the philosophical discourse is a conventional change of language; it is not the discovery of a fact about the world or the meaning of the word ‘responsible’.

5 The relation between science and artificial language philosophy of science

Whether traditional or ordinary language philosophy, partially naturalized or not, leads to fruitful interactions of science and philosophy of science is far from certain. It is, for example, not obvious how insights into the use of a term in ordinary language relate to scientific insights. Even when the methodology of ordinary language philosophy is applied to scientific language, it can discover inconsistent usage (Philipse 2009, §3) but cannot fix it without threatening to collapse into artificial language philosophy. And traditional philosophy has to establish its own access to facts about the world, besides the scientific route. Williamson (2007, §§6, 8) and Papineau (2009, §IV) consider thought experiments and thus ultimately intuitions to provide this access, but they both rely on contentious claims about the workings of the human mind. In the following, I will argue that the relation between science and artificial language philosophy of science is unstrained.

In a helpful overview, Hansson (2008) describes the ways in which philosophy relates to scientific disciplines. As they stand, however, these descriptions convey only sociological observations about the behavior of philosophers. Even if they could be explained on psychological grounds, a *justification* of the relations has to rely on some feature of philosophy itself. I will argue that the relations are justified and clarified when the methodology of artificial language philosophy is assumed.

New empirical results provide material for philosophical investigation. Hansson (2008, 477) describes a host of influences of scientific disciplines on philosophical work, but his examples mix concept formation, formal proofs, and empirical results. With respect to the empirical results, Hansson notes the influence of quantum mechanics and evolutionary biology on philosophy, the influence of psychology and neuroscience on the philosophy of mind, and the influence of linguistics on the philosophy of language. These examples show that some philosophical concepts (in the philosophy of mind, philosophy of language, etc.) are chosen to accommodate empirical results, and thus have to change when new results are found. This is a trivial implication of the way languages are chosen in philosophy as well as in the sciences.

New formal results provide material for philosophical investigation. Hansson (2008,

477) furthermore states that results in game and decision theory have provided moral philosophy with new problems for ethical analysis. Such results are purely formal in that they can be proved as theorems within the language. Because they are not empirical, they can also fall within the domain of philosophy, or they can suggest new language choices in philosophy, for example by establishing novel relations between concepts.

New concepts providing material for philosophical investigation. Hansson (2008, 477) further notes that game and decision theory have also provided new formulations of old problems in moral philosophy. While moral philosophy is outside the scope of this article, it seems clear that, to use Hansson's examples, psychology, neurosciences, linguistics, quantum mechanics, and biology have all led to new concepts. Carnap (1966, 187–189) automatically considers such conceptual work philosophical because it does not involve asserting or testing observational claims. But even with a more restrictive view of philosophy, it is clear that some philosophical concepts rely on scientific ones (by way of conditions of adequacy, for example), and therefore must change whenever the scientific concepts do. Since scientific concepts can change because of new empirical results, the relation between scientific and philosophical language choice provides yet another way for philosophical concepts to change in light of empirical results.

Methods and issues of philosophy are taken up by other sciences. According to Hansson (2008, 477), some issues and methods of philosophy have been taken up within other disciplines, and he gives the example of computer science, where the structures of concepts and thought processes have been investigated. Since conceptual and formal work is undertaken not only in philosophy but also in the sciences, it is unsurprising that science can join philosophy in developing concepts and establishing theorems. This has happened in the case of formal logic and mathematics, and to a great extent also in the empirical sciences. That currently computer science seems to be especially close to philosophical research may stem from the computer scientists' need for new languages that capture the structure of concepts and thought processes.

Philosophy is part of the community of interdependent disciplines. Hansson (2008, §3) notes the growing number of interdisciplinary endeavors and concludes on historical grounds that philosophy is part of the "community of interdependent disciplines". Successful philosophical investigations into natural or social phenomena, he claims, have always relied on results from other disciplines (such as the reliance of the philosophy of space and time on relativity theory).

Given the discussion so far, the interdisciplinary nature of philosophical research is obvious. The large role of language choice in the sciences is probably most evident in space-time physics: Relativity theory does predict new empirical phenomena, but it is

also a suggestion for a new language to accommodate old phenomena in a different way. Philosophers of space and time have had to evaluate this suggestion, and indeed have accepted its superiority over previous languages in many contexts.

Problems addressed experimentally or accurately become non-philosophical. Hansson (2008, 476f) also reviews the claim that many philosophical topics move into a dedicated field of science once clear answers are at hand. He gives the example of psychology, which parted from philosophy after the introduction of experiments.

Since accuracy in the rules of the application of a term is one desideratum of an explication, it is unsurprising that the fulfillment of this desideratum often marks the end of philosophical work. Conversely, empirical research involving auxiliary terms is only possible if there are meaning postulates to connect auxiliary and basic terms. The split between philosophy and psychology can therefore be seen as the result of the development of a precise language establishing such connections. Of course, once experimental research is pursued, it is still possible to change the language according to experimental results. A precise language can also allow formal work, so that Hansson's point also applies to, for example, mathematics and logic.

The autonomy of applied philosophy. According to Hansson (2008, §8), the philosophy of science is not an application of epistemology in the way that applied mathematics is an application of pure mathematics. Instead, he claims, philosophers of science develop their own theories which are related to—but not derivable from—epistemology.

Epistemology is usually more general than philosophy of science, aiming at the explication of concepts such as *belief*, *justification*, etc. (and more generally the formation of concepts) for as many contexts as possible. Philosophy of science, on the other hand, explicates concepts within the context of scientific theories and scientific practice. Given the different domains, it is to be expected that the explicata differ: Not all contexts that are relevant in epistemology are relevant in the philosophy of science, and some contexts that are very important in the philosophy of science only play a minor role in general epistemology.

Contrary to Hansson's suggestion, there is an analogy between the philosophy of science and the application of mathematics, namely when new mathematical concepts are developed in view of a specific application. A famous example is Dirac's " δ -function", which in fact cannot be treated as a function and which led to the development of the theory of distributions. The perception of a disanalogy between applied mathematics and applied epistemology may rest on a failure to distinguish between formal work and concept development: If a concept applies to a great variety of contexts, then the formal results for this concept will apply to each specific context, but there is no reason to assume that the best concept for a great variety of contexts is also the best concept for each

specific one.

Philosophical truths are not eternal. Hansson (2008, §6) suggests that many philosophers see their discipline as independent of empirical, synthetic results, which have no relevance in the philosopher’s realm of eternal, analytic truths. But, Hansson contends, Quine (1963) has shown that there is no uncontroversial line between analytic and synthetic statements, and philosophers who ignore empirical results (e. g., relativity theory) to arrive at claims that are “analytically true” (e. g., about time) are often just “demonstrably wrong”.

It is understandable that Hansson considers the analytic-synthetic distinction to be a problem for the connection between empirical science and philosophy, since analytically true sentences cannot be demonstrably wrong if such a demonstration is meant to be empirical. This is because the very definition of an analytic truth (as discussed here) is that it has no empirical import. The thorough critiques by Mates (1951), Martin (1952), and Kemeny (1963) of Quine’s attack on the analytic-synthetic distinction may thus also seem like a problem. However, there is no need to question the distinction to establish the relevance of empirical results for philosophical work. This is because even sentences that cannot be shown to be true or false by empirical research can be *chosen* to be true (or false). This amounts to concept formation, which, as discussed, can be and typically is heavily influenced by empirical results.

Hansson’s claim can be reinterpreted to state that a language chosen without reliance on empirical results is very unlikely to be able to accommodate them better than a language that was chosen with these results in mind. In this reading, Hansson’s claim is almost trivial. It also suggests that those philosophers who, for instance, do not consider relativity theory when explicating ‘time’ do not intend their explicatum to accommodate all the empirical results that relativity theory is meant to accommodate.

Philosophy of or with a discipline. With respect to philosophical endeavors related to the sciences, Hansson (2008, §7) distinguishes between the philosophy *of* science and philosophy *with* science. He states that philosophers of economics, for example, use the “tools of philosophy” to investigate how economists reason, so that philosophers relate to economists very much like social scientists to their objects of study. Philosophy with economists, on the other hand, consists of research conducted in collaboration with economist (for example, on the development of new representations of human beliefs, preferences, and norms).

Clearly, philosophers and scientists can work together when developing the same concepts and establishing formal proofs involving them. This is Hansson’s “philosophy with the sciences”, and it is to be expected in areas where the foundational concepts have not yet been developed fully. Then the construction of new representations is a major element

of research, and the formal and empirical elements cannot yet be pursued independently. Of course, it is also possible to suggest improvements of concepts already explicated in the sciences themselves and used with much success. Such improvements may, for example, be justified when dealing with problems that are not at the center of scientific research and that may not have carried much weight in previous explications.

Hansson’s description of “philosophy of science” is problematic because he does not specify “the tools of philosophy”. Given the comparison to social science, Hansson probably has the non-normative part of ordinary language philosophy in mind, since its application would lead to descriptions of the rules of language use in the sciences.

* * *

A closer look at the “tools” of artificial language philosophy, i. e. formal proof and language choice, clarifies the possibilities for a “philosophy of science” in Hansson’s terminology, and also reveals the possible relations between science and philosophy of science that are missing from Hansson’s list. Within artificial language philosophy, Hansson’s philosophy of science is probably best captured as the explication of concepts that are not explicated in the sciences themselves, but are still used in those contexts, including such general scientific concepts like *explanation* and *probability*. The concepts may not connect very well to others, and indeed may be confusing in certain contexts. Take the fine-tuning problem in physics. On the commonly used scales, the problem goes, the range of values in which the universal constants of physics allow life to exist are so small that the existence of life is very improbable, and therefore life is in need of an explanation (cf. Ratzsch 2009, §4.1). The actual usage of ‘explanation’, ‘probability’, and ‘scale’ in the sciences probably does not allow this inference. Therefore, the first step towards a solution of this problem is the explication of the terms, and this explication would fall in the domain of philosophy.⁸ If the explicata do not allow the inference, the explication of the terms is also the last step.

An example of a discipline-specific concept that, once introduced, has been explicated more extensively in philosophy than in the respective science itself is the notion of *gene* (cf. Waters 1994). In this case, however, Hansson’s distinction between philosophy *of* and *with* science becomes very blurry indeed. The explication of a discipline-specific concept is farther removed from Hansson’s philosophy with science when its goals are different from those in the respective discipline, leading to different conditions of adequacy. One instance of this is the explication of an initially discipline-specific concept for the simultaneous use in other disciplines. *Life*, for example, is comparably well-explicated within biology, but not for simultaneous use in robotics. A philosophical inquiry into the implications of artificial life may therefore have to develop its own explicatum. *Life* may also have to be explicated differently when used in ethical theories, and thus sometimes a scientific

⁸Note that it may be necessary to explicate the concepts differently for different scientific domains.

concept may have to be explicated for simultaneous use in a non-scientific domain. Waters (2004, §6) discusses the conditions of adequacy for such interdisciplinary explications in depth. Interdisciplinary explications are also desirable if the same term is already used for two slightly different concepts in two different fields, which can lead to fruitful interaction in some circumstances, but also abject confusion in others. ‘Information’ is a paradigmatic example.

Finally, there are concepts that are not used in science at all, but whose explication must take scientific results into account. *Personal identity* or *free will* may not occur (centrally) in the scientific literature, but for many contexts, their explication will have to take into account scientific results about, for example, the functioning of the brain and the predictability of individual behavior.

6 Conclusion

In artificial language philosophy, both science and philosophy of science engage in concept formation and, more generally, language choice, with often (but not always) differing conditions of adequacy and differing pragmatic goals. Accordingly, science and philosophy of science can interact and benefit from each other in a variety of ways. It is important to see that in artificial language philosophy of science these interactions and benefits are not simply postulated, but can be supported on methodological grounds.

It is also important to see that thanks to the possibility of reinterpreting many philosophical practices, artificial language philosophy can use many of the philosophical results achieved in ostensibly competing philosophical methodologies. Specifically, investigations of scientific language in ordinary language philosophy (including naturalized and experimental philosophy) can be interpreted as an identification of the rules of the language that scientists have developed to accommodate their experimental results. Investigations of traditional philosophy can be interpreted as a pragmatic search for a helpful language rather than for facts about the world.

It may be that any philosophical methodology has to be applicable to itself—specifically, one may take the view that if artificial language philosophy is to be used as a methodology for the philosophy of science, then it also has to be usable for the discussion *about* philosophy of science. In this case, I think the above shows that artificial language philosophy of science fulfills two important conditions of adequacy: It makes sense of both philosophical practice and the relation between science and philosophy of science. It can therefore be considered a good language choice.

References

- Andreas, H. (2010). Semantic holism in scientific language. *Philosophy of Science*, 77(4):524–543.
- Bergmann, G. (1949). Two criteria for an ideal language. *Philosophy of Science*, 16(1):71–74.
- Bergmann, G. (1957). The revolt against logical atomism—I. *The Philosophical Quarterly*, 7(29):323–339.
- Carnap, R. (1934a). *Logische Syntax der Sprache*, volume 8 of *Schriften zur wissenschaftlichen Weltauffassung*. Springer-Verlag, Wien. References are to the corrected translation (Carnap 1967).
- Carnap, R. (1934b). On the character of philosophic problems. *Philosophy of Science*, 1(1):5–19.
- Carnap, R. (1939). *Foundations of Logic and Mathematics*, volume I,3 of *Foundations of the Unity of Science. Toward an International Encyclopedia of Unified Science*. University of Chicago Press, Chicago and London. References are to the two-volume edition.
- Carnap, R. (1950a). Empiricism, semantics, and ontology. *Revue Internationale de Philosophie*, 4:20–40. References are to the slightly modified reprint (Carnap 1956, appendix A).
- Carnap, R. (1950b). *Logical Foundations of Probability*. University of Chicago Press, Chicago. References are to the 2nd edition (Carnap 1962).
- Carnap, R. (1952). Meaning postulates. *Philosophical Studies*, 3(5):65–73.
- Carnap, R. (1955). Meaning and synonymy in natural languages. *Philosophical Studies*, 6(3):33–47.
- Carnap, R. (1956). *Meaning and Necessity*. University of Chicago Press, Chicago, 2nd edition.
- Carnap, R. (1962). *Logical Foundations of Probability*. University of Chicago Press, Chicago, 2nd edition.
- Carnap, R. (1963). Replies and systematic expositions. In Schilpp, P. A., editor, *The Philosophy of Rudolf Carnap*, volume 11 of *The Library of Living Philosophers*, pages 859–1016. Open Court Publishing Company, Chicago and LaSalle, IL.

- Carnap, R. (1966). *Philosophical Foundations of Physics: An Introduction to the Philosophy of Science*. Basic Books, Inc., New York and London. Edited by Martin Gardner.
- Carnap, R. (1967). *The Logical Syntax of Language*. Routledge & Kegan Paul Ltd, London. Reprinted with corrections. Translation by Amethe Smeaton (Countess von Zeppelin).
- Chang, H. (2005). A case for old-fashioned observability, and a reconstructed constructive empiricism. *Philosophy of Science*, 72(5):876–887.
- Cohnitz, D. and Häggqvist, S. (2009). The role of intuitions in philosophy. *Studia Philosophica Estonica*, 2.2:1–14. Special issue: *The Role of Intuitions in Philosophical Methodology*, edited by Daniel Cohnitz and Sören Häggqvist. <http://www.spe.ut.ee/ojs-2.2.2/index.php/spe/article/view/84/53>.
- Daniels, N. (2008). Reflective equilibrium. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. The Metaphysics Research Lab, Center for the Study of Language and Information, Stanford University, Stanford, fall 2008 edition.
- Demopoulos, W. (2008). Some remarks on the bearing of model theory on the theory of theories. *Synthese*, 164(3):359–383.
- Dorr, C. (2010). Review of Ladyman and Ross (2007). In Gutting, G. and Gutting, A. F., editors, *Notre Dame Philosophical Reviews*. Philosophy Department, University of Notre Dame, Notre Dame, IN.
- Eshleman, A. (2004). Moral responsibility. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. The Metaphysics Research Lab, Center for the Study of Language and Information, Stanford University, Stanford, fall 2004 edition.
- Fedyk, M. (2009). Philosophical intuitions. *Studia Philosophica Estonica*, 2.2:54–80. Special issue: *The Role of Intuitions in Philosophical Methodology*, edited by Daniel Cohnitz and Sören Häggqvist. <http://www.spe.ut.ee/ojs-2.2.2/index.php/spe/article/view/57>.
- Feigl, H. (1958). Critique of intuition according to scientific empiricism. *Philosophy East and West*, 8(1/2):1–16.
- Feigl, H. (1969). De principiis non est disputandum? In Hospers, J., editor, *Readings in Philosophical Analysis*, pages 111–135. Routledge & Kegan Paul, London. Written in 1948.
- Feldman, R. (2008). Naturalized epistemology. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. The Metaphysics Research Lab, Center for the Study of Language and Information, Stanford University, Stanford, fall 2008 edition.

- Giere, R. N. (1985). Philosophy of science naturalized. *Philosophy of Science*, 52(3):331–356.
- Hansson, S. O. (2008). Philosophy and other disciplines. *Metaphilosophy*, 39:472–483.
- Hare, R. M. (1960). Philosophical discoveries. *Mind*, 69(274):145–162.
- Hempel, C. G. (1952). *Fundamentals of Concept Formation in Empirical Sciences*, volume II,7 of *Foundations of the Unity of Science. Toward an International Encyclopedia of Unified Science*. University of Chicago Press, Chicago and London. References are to the two-volume edition.
- Kauppinen, A. (2007). The rise and fall of experimental philosophy. *Philosophical Explorations*, 10(2):95–118.
- Kemeny, J. G. (1963). Analyticity versus fuzziness. *Synthese*, 15(1):57–80.
- Kim, J. (1988). What is “naturalized epistemology?”. *Philosophical Perspectives*, 2:381–405.
- Kuipers, T. A. F. (2007). Introduction. Explication in philosophy of science. In Kuipers, T. A. F., editor, *General Philosophy of Science—Focal Issues*, volume 1 of *Handbook of the Philosophy of Science*, pages vii–xxiii. Elsevier, Amsterdam.
- Kyburg, Jr., H. E. (1990). Theories as mere conventions. In Savage, C. W., editor, *Scientific Theories*, volume 14 of *Minnesota Studies in the Philosophy of Science*, pages 158–174. University of Minnesota Press, Minneapolis, MN.
- Ladyman, J. and Ross, D. (2007). *Every Thing Must Go: Metaphysics Naturalized*. Oxford University Press, Oxford. With David Spurrett and John Collier.
- Lewis, D. (1970). How to define theoretical terms. *The Journal of Philosophy*, 67(13):427–446.
- Lutz, S. (2009). Ideal language philosophy and experiments on intuitions. *Studia Philosophica Estonica*, 2.2:117–139. Special issue: *The Role of Intuitions in Philosophical Methodology*, edited by Daniel Cohnitz and Sören Häggqvist. <http://www.spe.ut.ee/ojs-2.2.2/index.php/spe/article/view/65>.
- Martin, R. M. (1952). On ‘analytic’. *Philosophical Studies*, 3:42–47.
- Mates, B. (1951). Analytic sentences. *Philosophical Review*, 60:525–534.
- Mates, B. (1958). On the verification of statements about ordinary language. *Inquiry*, 1:161–171.

- Maxwell, G. and Feigl, H. (1961). Why ordinary language needs reforming. *The Journal of Philosophy*, 58(18):488–498.
- Nadelhoffer, T. and Nahmias, E. (2007). The past and future of experimental philosophy. *Philosophical Explorations*, 10(2):124–149.
- Nichols, S. and Knobe, J. (2007). Moral responsibility and determinism: The cognitive science of folk intuitions. *Noûs*, 41:663–685.
- Papineau, D. (2009). The poverty of analysis. *Aristotelian Society Supplementary Volume*, 83(1):1–30.
- Philipse, H. (2009). Can philosophy be a rigorous science? *Royal Institute of Philosophy Supplement*, 84:155–176.
- Przełęcki, M. (1969). *The Logic of Empirical Theories*. Monographs in Modern Logic Series. Routledge & Kegan Paul/Humanities Press, London/New York.
- Quine, W. V. O. (1963). Two dogmas of empiricism. In *From a Logical Point of View. Logico-Philosophical Essays*, pages 20–46. Harper & Row, New York et al., 2nd edition.
- Quine, W. V. O. (1969). *Ontological Relativity and Other Essays*. Number 1 in The John Dewey Essays in Philosophy. Columbia University Press, New York and London.
- Ratzsch, D. (2009). Teleological arguments for God’s existence. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. The Metaphysics Research Lab, Center for the Study of Language and Information, Stanford University, Stanford, winter 2009 edition.
- Reichenbach, H. (1938). *Experience and Prediction; an Analysis of the Foundations and the Structure of Knowledge*. University of Chicago Press, Chicago.
- Reichenbach, H. (1951). The verifiability theory of meaning. *Proceedings of the American Academy of Arts and Sciences*, 80(1):46–60.
- Rorty, R. (1967a). Introduction: Metaphilosophical difficulties of linguistic philosophy. In Rorty (1967b), pages 1–39.
- Rorty, R., editor (1967b). *The Linguistic Turn: Recent Essays in Philosophical Method*. University of Chicago Press, Chicago and London.
- Sosa, E. (2007). Experimental philosophy and philosophical intuition. *Philosophical Studies*, 132(1):99–107.
- Stotz, K., Griffiths, P. E., and Knight, R. (2004). How scientists conceptualize genes: An empirical study. *Studies in History and Philosophy of Biological and Biomedical Sciences*, 35(4):647–673.

- Suppe, F. (1971). On partial interpretation. *The Journal of Philosophy*, 68(3):57–76.
- Sytsma, J. (2010). The proper province of philosophy: Conceptual analysis and empirical investigation. *Review of Philosophy and Psychology*, 1(3):427–445. Special issue: *Psychology and Experimental Philosophy (Part II)*, edited by E. Machery, T. Lombrozo, and J. Knobe.
- Waters, C. K. (1994). Genes made molecular. *Philosophy of Science*, 61(2):163–185.
- Waters, C. K. (2004). What concept analysis in philosophy of science should be (and why competing philosophical analyses of gene concepts cannot be tested by polling scientists). *History and Philosophy of the Life Sciences*, 26(1):29–58.
- Williamson, T. (2007). *The Philosophy of Philosophy*, volume 2 of *The Blackwell/Brown Lectures in Philosophy*. Blackwell Publishing, Malden, MA.