Varieties of Conceptual Analysis

Max Kölbel, University of Vienna

Abstract: What exactly does conceptual analysis consist in? Is it empirical or a priori? How does it support philosophical theses, and what kinds of thesis are these? There is no consensus on these questions in contemporary philosophy. This paper aims to defend conceptual analysis by showing that it comprises a number of different methods and by explaining their importance in philosophy. After setting out an initial dilemma for conceptual analysis, the paper outlines a minimal ecumenical account of concepts, as well as an account of concept possession and concept employment. On the basis of these accounts, the paper then argues that there are both empirical and a priori forms of conceptual analysis, and that each can be defended as legitimate methods. The philosophical interest of conceptual analysis, however, resides in relying on all three types of method in the service of answering philosophical concerns. This is illustrated by three sample cases.

1. A Dilemma for Conceptual Analysis

Conceptual analysis has seemed to many philosophers to be an important, if not central philosophical activity.[1] Philosophers, on such a view, analyse concepts like the concept of knowledge, of free will, justice, personhood, identity, truth, beauty, reason, obligation and many more.[2] Conceptual analysis has had the reputation of being a method, that can be carried out in an *a priori* manner, and from the proverbial armchair, by consideration of hypothetical cases (fictional scenarios, thought experiments). However, conceptual analysis faces the following dilemma. Concepts, the conceptual

---

[1] For example to Ryle (1949), Carnap (1950), Grice (1958, 1987), Strawson (1963, 1992), Jackson (1998), Chalmers & Jackson (2001) or Nolan (2009).

[2] They may *also* theorize about knowledge, free will, justice etc. as opposed to the concepts thereof. It is here assumed that these are *distinct* objects of investigation. This essay is about conceptual analysis and tries to explain the interest of various forms of it. It is *not* taken for granted that the purpose of analysis of the concept of, e.g., justice, is the acquisition of knowledge about justice.

analyst's objects of research, can either (i) be identified as the concepts particular individuals or groups actually employ in thought. Alternatively, (ii) concepts can be identified by specifying their defining characteristics (such as rules of correct application, inference rules, possession conditions, etc). Each of these two options seems problematic.

On the first horn of the dilemma the object of investigation is empirical: actual episodes of thought, psychological processes or, in so far as concepts are shared, group interactions, etc. This leads to at least two problems: on the one hand, it is not clear how armchair consideration of hypothetical cases can provide adequate evidence on this empirical subject matter. (Experimental philosophers tend to point out shortcomings of this sort.) Secondly, it seems problematic to say that philosophers are centrally pursuing empirical questions on psychological or sociological terrain—not just because they would be dabbling and poaching in foreign academic disciplines, but because philosophy would seem deprived of its own proper subject matter.[3]

On the second horn of the dilemma, the subject matter, i.e. concepts as abstract objects, identified by definition, may be susceptible to purely a priori methods, but there is doubt about the relevance of such a priori investigations to philosophical concerns. In Strawson's words (commenting on Carnap):

> [T]ypical philosophical problems about the concepts used in non-scientific discourse cannot be solved by laying down the rules of use of exact and fruitful concepts in science. To do this last is not to solve the typical philosophical problem, but to change the subject. (Strawson 1963, p. 505)

The worry is that if one identifies a concept by laying down a definition it may be possible to arrive at purely a priori conclusions about the concept thus defined, but there is no guarantee that these conclusions concern the original philosophical worry about, e.g., justice or freedom. They may simply be irrelevant.

What lessons should be drawn for philosophy and the project of conceptual analysis? Leaving aside the large faction of those who carry on unreflectively without acknowledging any problem, contemporary responses can be classified as follows: a

---

[3] This will seem a problem both to philosophers who, like Bealer 1998, regard philosophy as distinct from empirical science, and to those, like Papineau 2009, 2014 or Kornblith 2002, who think that philosophy is continuous with science but still has its proper subject matter.

rejection of conceptual analysis, an endorsement of the first horn and an endorsement of the second horn.

First, some reject conceptual analysis altogether, i.e. they deny that conceptual analysis is an important philosophical activity. These philosophers will interpret what might look like the examination of, say, concepts of knowledge, free will, etc via consideration of fictional cases as in fact a method for investigating knowledge, free will etc themselves.

Secondly, some embrace the first horn of the dilemma and insist that conceptual analysis targets psychological and/or sociological facts about humans. This can involve the recommendation of proper experimental methods (usually methods known from social science or psychology) to investigate this subject matter. It can also involve defending traditional armchair methods to some extent, not as a priori, but as empirical methods. This response is common among contemporary experimental philosophers, but it has been widespread already before the recent movement in experimental philosophy.

A variant of this second approach, not always distinguished properly, involves regarding as the target of analysis not concepts but meanings of natural language expressions. On this view, the point of considering fictional cases and thought-experiments is to gather data about language use and thereby about the semantic properties of expressions.[4]

Thirdly, there are those philosophers who embrace the second horn and accept that conceptual analysis uses purely a priori methods to study abstract entities (concepts) that are identified by definition.

The first type of response has been the dominant attitude for some time: since Quine's criticism of analyticity, a certain hostility towards conceptual analysis has been prevalent. There have been serious attempts to explain how the method of considering fictional cases can yield knowledge about free will, knowledge, and other philosophical targets of analysis without a detour via concepts or meanings (to mention just a few examples: Yablo 1993, Kornblith 2002, Williamson 2007, 2018; Cappelen 2012; Levin

---

[4] The variant could also be seen as a variant of the first response: rejecting conceptual analysis (both horns) and reinterpreting what philosophers do—in this case as meaning analysis. Very often, philosophers who see themselves as doing conceptual analysis take for granted that this involves the analysis of the meaning of linguistic expressions that express the concepts in question. Grice 1958, Strawson 1992, Jackson 1998, Thomasson 2015 are typical examples.

2014; Papineau 2014; Bengson 2015; Horvath 2018). The second type of response has also had its supporters, e.g. Grice 1958, 1987; Strawson 1963, 1992; Jackson 1998; Chalmers & Jackson 2001; Chalmers 2012; Machery 2017). Many have long regarded the point of considering hypothetical cases as eliciting judgements that reveal something about our concepts—or on the mentioned variant: about our linguistic competence (e.g. Grice 1958, 1987; Strawson 1992, Chalmers 1998, Nolan 2009 and Thomasson 2015)[5]. It is this conception that is the main target of criticism from experimental philosophers. They take for granted that conceptual analysis targets psychological (or perhaps psychological-cum-sociological) entities and criticize armchair methods as methods for targeting these empirical phenomena. What they propose is primarily an improvement, or even replacement, of traditional philosophical methods for pursuing conceptual analysis (see e.g. Machery 2017, see Sytsma & Buckwalter 2016 for an overview).

The third type of response tends to be neglected or forgotten in this context—even though it is the only approach on which conceptual analysis has a chance of being a genuinely *a priori* pursuit that does not postulate extravagant philosophical faculties.

It is my aim in this paper to put forward a pluralistic view of conceptual analysis that not only recognizes the importance of both empirical and a priori methods of conceptual analysis, but also stresses that they complement one another in the pursuit of further philosophical goals. There is a variety of questions about concepts that philosophers pursue, and there is a corresponding variety of methods of conceptual analysis. Often, finding out about our concepts has both an empirical and an a priori aspect. This becomes especially clear when philosophers pursue questions about concepts not just with a descriptive, but with a prescriptive or normative agenda. The view proposed here neither precludes nor defends the possibility of there also being worthwhile philosophical questions that do not involve analysing concepts (as proponents of the first response often claim).

---

[5] I will use the term "judgement" when some theorists use "intuition". The term "intuition" is unhelpful in the current methodological discussion, because it tends to be understood as already connoting a certain justificatory role (see, e.g. Kauppinen 2013). In our context, it will be clearer to speak of judgements with certain justificatorily relevant characteristics, such as being experientially justified, spontaneous, made by an authoritative or competent thinker, etc. Sometimes, intuitions are thought of not as judgements but as experiences or other quasi-perceptual states (see, e.g. Bengson 2015). Here, too, I try to keep things clear by not using the term "intuition". In any case, those theorists using the term "intuition" are not usually using it to denote a source of evidence in a project of conceptual analysis.

There is a lot of controversy about concepts and no consensus on what concepts are. These controversies are mostly orthogonal to the methodological proposals concerning conceptual analysis that I will be making. So I shall begin in §2 by developing a minimal account of concepts that will allow me to present my pluralistic methodology of conceptual analysis in such a way that it can be articulated within all approaches to concepts, thus making my proposal relevant independently of the view one may hold about concepts. This will allow me in §3 to draw a straightforward distinction between empirical and a priori questions about concepts. In order to give concrete substance to some of my proposals, I will need to add some flesh to the minimal skeleton account of concepts, which I will do by sketching a view of concept employment and its behavioural manifestations in §4. Equipped with these accounts of concepts and concept employment, I shall then proceed in §§5–7 to explain some a priori and empirical methods by which concepts can be examined, and how these can complement one another in the pursuit of philosophical aims.

2. A Minimal Account of Concepts

What are concepts? Some basic ingredients seem to be present in almost all accounts of concepts: concepts are, in some sense, constituents or tools of thought. They classify in that they (typically) apply selectively to some and not to other things.[6] For example, the concept of a tree, or better *a* concept of tree, is *employed* by a thinker who believes or thinks that the recent storm has uprooted some trees. One needs, in some sense, to *have* or *possess* such a concept, in order to be able to form that belief or to think that thought. Such a concept in some sense *occurs* in, or *partially constitutes* the belief. It is *correct* to apply the concept to some things, but not to others: in order for the belief that the storm has uprooted some trees to be true, it will not be sufficient that an electricity pylon get uprooted. It needs to be something to which the *tree* concept in question can be correctly applied.

This much is perhaps uncontroversial common ground in discussions about concepts. In sum: (i) concepts can be *possessed* by thinkers; (ii) *employed* by those

---

[6] My talk of concepts in general "applying selectively" should not be taken to rule out that some concepts apply to everything, and some to nothing at all.

thinkers in episodes of thought (the concepts can be said to *occur* in such thoughts); and (iii) it can be *correct* or not to apply a concept to an object.[7]

But there are many different conceptions of concepts that interpret this minimal common ground.[8] Usually, the different conceptions are assumed to be rivals, and to be incompatible with one another. One important group will think of concepts as abstract objects that are constituents of another type of abstract object, namely of propositions or Fregean thoughts (for example Peacocke 1992, Zalta 2001 or Båve 2019). Members of this group will have some view about how individual thinkers are related to concepts and propositions to make sense of the idea of them "employing" or "having" concepts. Another important group think of concepts as concrete mental representations in the minds of thinkers (e.g. Millikan 2000, Margolis & Laurence 2007, Carey 2009). According to Machery 2009, this psychological conception views concepts as "bodies of knowledge" of a certain kind that are "used by default in the processes underlying higher cognition" (Machery 2009, p. 50). Members of this group will have some view on whether and how different thinkers can be employing the same concept, and in what sense some concept applications are correct while others are not.

I do not share the assumption that rivalling conceptions of concepts are generally incompatible with one another.[9] Adopting one of these conceptions need not put one in any substantial (non-verbal) conflict with any of the others—although, of course, there is a general obligation to make clear what one is talking about. In this section, I shall distinguish some basic conceptions of concepts and how these different conceptions can be employed to talk in different but compatible ways about the same phenomena. The point of this is to identify some points at which there is no non-verbal disagreement between the conceptions, in order to eliminate some sources of misunderstanding.

One of the basic points on which different conceptions of concepts diverge is whether concepts are to be unchangeable entities or whether they can change over time, perhaps coming into existence, or ceasing to exist, at some point in time.

---

[7] NB: element (iii) of this minimal common ground says merely that in principle it can be correct or incorrect to apply a concept in a given case. This says nothing about how correctness of application is determined, whether correctness is defined for all possible cases (rather than only partially), or whether it is possible to articulate such conditions, e.g. by formulating necessary and sufficient conditions. Thus it is a truly minimal assumption.

[8] See for example Margolis and Laurence 2014, for an overview.

[9] Thus, I agree with Machery 2009 that what Machery calls "psychological" and "philosophical" conceptions of concepts are meant to address very different theoretical questions. I also agree with Glock 2011 and Camp 2015 that the conceptions can complement one another.

Studying the history or development of concepts seems to be an endeavour that calls for the latter type of conception. If, for example, I want to discuss the history of the concept of a smartphone it looks like it would be convenient to think of this concept as a concept that was created at some point, perhaps in the late 1990s, and that has possibly undergone various changes since then. Similarly, one may legitimately wonder whether the (a certain group's) concept of plagiarism has changed since the advent of the internet. Of course, we must distinguish the question whether smartphones and plagiarism have changed from the question whether *the concepts* of smartphone and of plagiarism have changed. For it is obvious that plagiarism has changed since the advent of the internet (a new type of source has become available), while it is not so obvious whether the concept of plagiarism has changed. One reason to say that a concept has changed would be if the conditions for correct application of the concept change. Arguably, the concept of a planet changed in this way when the condition of being gravitationally dominant in one's neighbourhood became one of the conditions an object needs to meet in order for it to be correct to apply the concept to it.

Clearly, if I want to describe matters in such a way that a concept comes into existence and changes, I must conceive of concepts as objects that persist over time and can undergo changes. Any such conception will involve (or presuppose) some view as to what constitutes concept identity over time, and perhaps also a view about how concepts come into existence. One common approach here is to individuate concepts by way of the linguistic expressions that are used to express judgements in which the concepts occur. Thus, one could say that the concept of plagiarism is simply the concept expressed by the word "plagiarism" in English (and perhaps its translations)[10]. Another approach is that of introducing a deference relation, so that one episode of concept employment counts as employing the same concept as another, if one stands in the deference relation to the other. The deference relation could involve one thinker's intention to use the same concept as another thinker (or the same thinker on another occasion).[11]

However, it is also possible to conceive of concepts as unchanging abstract entities. Studying the logical properties of concepts, and the logical relations among concepts,

---

[10] See, e.g. Jackson 1998, Thomasson, 2015. Of course, there are also further questions about what it is for a word to express a concept and also about the identity of words, or expressions, over time.
[11] Sainsbury & Tye 2011, 2012 pursue this type of approach.

seems to call for this type of conception. For example, in order to study whether a concept of truth is incoherent due to the liar paradox, it is convenient to focus on principles governing the correct application of the concept in question, in this case the equivalence principle. If I want to know whether a concept of freedom is such that an agent whose actions are causally determined could be free, it is again convenient to think of the concept of freedom in question as being governed by some principles concerning the conditions under which it can be correctly applied—in this case some principles about what it takes for an agent to be free. For these sorts of enquiry it is convenient to conceive of concepts as characterised constitutively by principles concerning correct application because the logical properties at which the inquiry aims depend on such principles. Thus it is no surprise that in this sort of project, philosophers typically start by articulating conditions of correct application or principles governing the concepts they are interested in.

A conception of concepts according to which concepts change over time—especially in so far as they change the principles of correct application to which they are subject—would be unnecessarily cumbersome if the focus of enquiry is precisely on those principles and their logical consequences. However, this is mere inconvenience. Of course I can frame my examination about a certain unchangeable concept of truth as an examination of the equivalence principle, a principle to which the changeable concept of truth used by a certain population is subject at a certain point in time. Or I could frame the examination simply as an examination of the equivalence principle and its consequences, not even using the term "concept". A historian of concepts using a conception of changeable concepts should have no objection at all to the other theorist's different conception of unchangeable concepts.

Could the historian of concepts also symmetrically dispense with his or her conception of changeable concepts? Again, it would be inconvenient, but possible. Rather than saying that some individual's or group's concept of a smartphone was created and subsequently underwent changes with regard to the principles that govern it, the historian could instead say that a certain concept user, or community of concept users, at one point used one concept, constitutively characterised by one set of principles, and at a later point used a distinct concept constitutively characterised by different principles.

There need therefore be no disagreement of substance between conceptions of concepts as changeable or as unchangeable. One and the same theorist could coherently and productively use both conceptions provided she made sure to avoid confusions— for example by introducing disambiguating terminology. For example, a theorist could say that concepts$_u$ are unchangeable and constitutively governed by certain principles regarding correct application, while concepts$_c$ can change their principles of correct application over time and are individuated by, for example, chains of deference or words used to express them. Such a theorist might even construe concepts$_c$ as constituted by time-slices in such a way that each time-slice is (or determines) a unique concept$_u$.

Another variation in conceptions of concepts worth discussing briefly is that between individual and social conceptions of concept. Amongst those who think of concepts as changeable continuants whose history or development can be examined, some mean to speak about (recurrent) representations in individual minds as manifested in individual behaviour (e.g. Millikan 2011); while others mean to speak about recurrent representations that are shared across different individuals, as manifested in social interactions (e.g. Sainsbury & Tye 2012). In my view, there need again be no substantial disagreement. Where the social conception speaks of cross-subject identity of concepts, the individualist will speak of mere cross-subject "classification" (see Millikan 2011). In other words, the socialist's concepts$_{cs}$ may be classes of the individualist's concepts$_{ci}$. Of course, whether the similarities between the concepts$_{ci}$ of distinct individuals are ever great enough to make for sufficiently interesting concepts$_{cs}$ that are classes of concepts$_{ci}$, may well be a substantial point of debate.

Related to social versus individualistic conceptions of concepts is the question of externalism versus internalism about concepts. According to externalism, it can depend in part on a thinker's external environment, which concept$_u$ she is employing, or what the conditions of correct application of their concept$_c$ are at a given point in time. Internalism denies this. What has been said so far should be compatible with externalism and externalism alike. (We will revisit externalism when discussing empirical conceptual analysis below.)

Let me return to the minimal common ground about concepts mentioned at the beginning: (i) concepts can be *possessed* by thinkers; (ii) thinkers who possess a concept can *employ* that concept in an episode of thought (the concept can *occur* in that

thought); and (iii) it can be *correct* or not to apply a concept to an object. We can now see that these three points are compatible with several different conceptions of concepts, and that these different conceptions themselves do not yet indicate any non-verbal disagreement between their proponents.

3. Empirical and A Priori Discovery

Given that minimally, concepts can occur in the thoughts of thinkers that possess them, and are also subject to some conditions of correct application, we can distinguish two ways of conducting research about them: empirical and a priori. Suppose a researcher stipulates that she is going to examine a concept C that is characterized by feature F. Then it is possible to examine the deductive consequences of C having F, and the results of this examination will be a priori, for they relied only on stipulation and deduction.[12] But it is also possible to conduct empirical research on concept C, as characterized, by considering empirical consequences of C having F. This second examination will be empirical, and not a priori.

For illustration, consider some examples. Suppose the conceptual analyst stipulates that she is going to examine the concept, call it "*true*1", expressed by a certain group G (during time-span t) by the expression "true". Starting with this stipulation, it will be a priori that when a member of G uses the word "true" (during t), she expresses the concept *true*1. It will be an a posteriori question whether *true*1 is subject to the equivalence principle, which states that any proposition that p falls under *true*1 iff p. Empirical data that could be taken to bear on the issue might be data concerning the behaviour of members of G, and the evidential significance of such data will depend on further hypotheses and possibly further stipulative assumptions. For example, the tendency (in favourable conditions) of members of G to accept instances of the schema 'It is true that p if and only if p.' might count as data in favour of *true*1 being subject to the equivalence principle.

Suppose, however, the conceptual analyst stipulates that she is going to examine a concept, call it "*true*2", whose conditions of correct application are governed by the equivalence principle. If these are the terms of the investigation, the question whether *true*2 is governed by the equivalence principle is not empirical. By contrast, the

---

[12] Thus, the conclusions drawn would be analytic in the minimal and inoffensive sense of deriving from stipulations and deduction alone (compare Frege 1884, Juhl & Loomis 2009, ch. 6).

question whether members of *G* (during *t*) express *true*2 by their utterances of "true" is empirical.

Before moving on I want to make three remarks about the distinction between a priori and empirical questions about concepts that I have introduced. First, the distinction classifies questions about concepts in unexpected ways. There is a tendency in the literature to assume that it is so-called psychological conceptions of concepts that give rise to empirical questions, and so-called philosophical or abstract conceptions of concepts that allegedly give rise to a priori questions. My distinction opens up both empirical and a priori questions on either type of conception. The crucial question, each time, is how the object of investigation is framed, what initial stipulations the theorist makes. Thus, there can be both a priori and empirical questions on each conception of concepts. In particular, it is not the case that philosophical conceptions of concepts as abstract objects automatically lead to a purely a priori enquiry.

Secondly, the distinction is not affected by the usual scepticism about a priori methods which derives from criticism of a certain notion of analyticity. While there are important philosophical issues about the epistemology of deduction, and while there may also be issues about what happens when a stipulation is made, these are not the issues that drive mainstream philosophical scepticism about analyticity. So I believe that it is legitimate, for current purposes, to assume that stipulation, or stipulative definition, together with deduction, can give rise to a priori justification and knowledge.[13]

Thirdly, and correlatively: to the extent to which the account of a priori questions here offered is innocuous and immune to standard criticisms, it may also seem to be far too modest to give rise to any interesting account of a priori conceptual analysis. What

---

[13] This sort of analyticity is not even a target of Williamson's recent objections: once it is stipulated that "every" is not to be existence-entailing, and after it has been stipulated for a range of common nouns '*F*' that '*F*' is either true or false for every object, claims of the form 'every *F* is an *F*' will be analytic even in Williamson's strong sense that competence entails assent (cf. Williamson 2007). Or, to mention another Williamsonian example: even Vann McGee will agree that a conditional connective that is stipulated to support modus ponens supports modus ponens (cf. Williamson 2003). Compare Juhl & Loomis 2009, ch. 6, who defend the epistemological significance of a minimal notion of analyticity*; as well as Soysal 2018 for a similar-spirited objection to Williamson's arguments.

    This sort of apriority is also not the target of recent criticisms of the Jackson-Chalmers approach, e.g. Stalnaker 2001, Laurence & Margolis 2003, Schroeter 2006 all criticize Chalmers and Jackson because they claim that we have a priori knowledge of what our concepts are (what our terms mean). Thus, their target is our second response to the intitial dilemma.

is interesting about the deductive consequences of stipulations? In his book on conceptual analysis, Jackson writes:

> If I say that what *I* mean —never mind what others mean—by a free action is one such that the agent would have done otherwise if he or she had chosen to, then the existence of free actions so conceived will be secured, and so will the compatibility of free action with determinism. … But … [I won't have] much of an audience. I have turned interesting philosophical debates into easy exercises in deductions from stipulative definitions together with accepted facts. (1998, p. 31)

Stipulative definitions and their deductive consequences, however, can be relevant and interesting if the definitions capture (aspects of) the concepts we actually use (or, as in Jackson's own approach, if they capture the meanings of the expressions we actually use). Thus, a priori results can be interesting for empirical reasons. Moreover, it is a mistake to think that deduction is always an easy exercise: many results in deductive logic are far from easy.

It will be the task of §§5–7 to fill these abstract remarks on a priori and empirical conceptual analysis with concrete content.

4. Employing Concepts

I have just made a general distinction between a priori and empirical questions about concepts merely on the basis of the minimal common ground about concepts: that concepts can be *employed* in thought by thinkers who *possess* them, and that they are subject to some *conditions of correct application*. Typically empirical conceptual analysis identifies a concept by stipulating that it is the concept employed in a given episode of thought or expressed by a given linguistic expression in a given speech community, and then proceeds to collect empirical data to draw conclusions about the correctness conditions of the concept thus identified. So, in order to give a more substantial characterisation of this empirical method, it will be necessary to flesh out the minimal common ground about concepts with a more substantial account of what is involved in concept possession and concept employment, and how they are manifested, so as to be able to explain what types of behaviour can be counted as empirical data about concepts identified in the way suggested.

What is it for a thinker to possess a concept, and, in a given episode of thought, to employ that concept? At a minimum, in order to possess a concept, a thinker needs to have a classificatory mechanism that is in some way sensitive to the concept's

correctness conditions (more on sensitivity shortly). "Classificatory mechanism" here simply means an ability to classify items that can be repeatedly exercised. For example, let us say that the concept *bike* is a concept that can be correctly applied to all and only bicycles. Then possession of the concept *bike* would involve a stable ability to classify objects in a way that is sensitive to these correctness conditions (i.e. sensitive to objects being bicycles) and to store corresponding information about objects. *Employing* a concept will then simply amount to exercising the ability that constitutes possession of the concept, and to do so in the formation of a belief.

I said that the possession of a concept requires a classificatory mechanism that is sensitive to the correctness conditions of the concept. I need to spell this out a little. A straightforward way for a classificatory mechanism to be sensitive to correctness conditions is for it simply to classify reliably in accordance with correctness conditions. This could mean that the mechanism reliably categorises an item only if it meets the correctness conditions (or if the thinker believes that it does). However, concept possession should not preclude regular mistakes, so this (i.e. reliably correct application) cannot be the only way to be sensitive. There should be room for unreliability among concept possessors. Thus, sensitivity to correctness conditions may be realised by a mere disposition to classify correctly under favourable conditions of classification. If conditions are rarely or never favourable, then sensitivity in this sense is perfectly compatible with frequent error, for frequent absence of favourable conditions can generate frequent incorrect application.

But there can also be concept possessors who would not even be reliably right if conditions were favourable. This can happen when sensitivity to the correctness conditions is socially mediated and takes the form of deference to other concept users: I may depend in my classifications on the testimony of authorities or experts, who in turn are sensitive in a more direct or autonomous way to correctness conditions (see Burge 1979). Thus, even though I am not fully competent regarding the application conditions of the concepts *arthritis*, say, I still count as using the same concept as the experts—and it is they whose classifications under favourable conditions determine the application conditions.

This brings me to the social character of concept possession and concept employment. Individuals may possess and employ concepts not only in virtue of their own dispositions to form beliefs, but also in virtue of membership in a group whose

members stand in social relations of, for example, authority, deference, etc. Individual concept users may be sensitive to the correctness conditions of a concept only in the sense that they stand to be corrected in their classifications by other users who are sensitive in a more direct way, for example via a disposition to correct application under favourable conditions. Moreover, concepts are usually acquired in a social way: fully competent thinkers serve as models for novices, novices emulate the models and thus acquire classificatory habits that are calibrated with their models.

In order to honour the ecumenical approach to concepts so far pursued, let me say that it is, of course, possible to pursue a completely individualistic view of concepts and their employment. On such a view, one may have to count the expert in a community as employing a distinct concept from the non-expert. When I, the non-expert, have a belief that I would express by saying "I have arthritis.", I employ a distinct concept from the expert who has a belief they would express by those words. Perhaps my concept is the concept that it is correct to apply to an ailment if and only if the medical experts call it "arthritis", while the expert's concept is one that it is correct to apply to an ailment iff it is an inflammation of a joint. How promising or useful this approach is need not be addressed here. For our purposes we just need to note such an individualistic approach to concepts and their employment will have a different view about what data about the manifestation of concept use show about the concepts employed. The view of conceptual analysis here defended is compatible with this individualistic view—even though it will change what type of empirical evidence is relevant and how it is interpreted.

To sum up: concept *possession* involves a classificatory mechanism, i.e. a stable ability to form beliefs of a certain sort that is sensitive to the concept's correctness conditions. This sensitivity can take various forms: for example, a disposition to classify in accordance with the correctness conditions under favourable conditions, or deference to the judgements of authoritative users of the concept who, perhaps, in their turn classify in accordance with the correctness conditions when conditions are favourable. Concept *employment* can then be viewed simply as the exercise of the corresponding classificatory mechanism.

This account of possession and employment provides an indication of how behaviour can indicate concept employment: classificatory mechanisms can be manifested when exercised (they are dispositions). In particular, if I want to answer the

empirical question whether a concept possessed by a given individual is subject to certain correctness conditions, then that individual's employments of the concept under favourable conditions will be relevant empirical data. If the question is one about a given concept shared by a group, then, depending on the precise social nature of the concept, the employments of the concept by the relevant authorities within that group may prove to be the most relevant data.

Two further clarificatory remarks about concepts, their possession and employment may be helpful: First, given my characterizations, there may well be several rival concepts possession and employment of which makes equally good sense of a person's manifested dispositions. In that case, it may be indeterminate whether in a given episode of thought, a person is employing a particular concept rather than one of the rivals. For our purposes, we can leave open whether this happens: the usefulness of conceptual analysis, as portrayed in this paper, is compatible with the relationships of possession and employment being vague or underdetermined. Secondly, this account of concepts, their possession and their employment is compatible with a kind of metaphysical minimalism about concepts: to describe an episode of thought or an interaction among several thinkers in terms of concepts employed is merely one way of describing that episode and those interactions in a way that systematizes certain regularities. Concepts, like propositions, could simply be "measures" of the mind (cf. Matthews 1994).

5. Empirical Investigations of Concepts

It may seem, then, that the only substantial questions about concepts are empirical questions regarding the concepts we actually employ in given episodes of thought, or the concepts we express when using certain linguistic expressions. For, if Jackson is right, easy deductions from stipulative definitions (by themselves) are irrelevant to the questions that interest philosophers. I do not believe that all purely a priori issues are easy or philosophically uninteresting, but that is the subject of the next section. In this section, I shall elaborate on empirical research about concepts in philosophy, starting with what seems to be a widespread method: the method of cases.

It is sometimes thought that philosophers can investigate a concept by consulting their own judgements [14] involving that concept. Often, philosophers will consider elaborate, even fantastical, scenarios and adjudicate whether a concept of interest applies or does not apply to a given entity in that scenario. Thus, an investigation into the concept of a person may proceed by looking at various fictional cases and drawing on judgements about who is a person, or counts as the same person, in these cases. Or an investigation about the concept of free will may proceed by looking at various scenarios in which an agent acts and draw on judgements as to whether the agent acted freely. This procedure is often described by saying that these philosophers are conducting "thought experiments".

To begin with, let me emphasise that I am treating this sort of method as empirical, not a priori. Since these thought experiments can be conducted from the proverbial armchair, they may seem superficially to be an a priori investigation. On my view, however, the procedure can be methodologically justified only on the assumption that it yields empirical evidence. How can it be justified?

Following the distinction made in §3, empirical investigations involve identifying a concept by making some stipulation about the concept one is interested in, and then to gather empirical evidence about the concept thus identified—using hypotheses about concept employment of the sort described in §4. For example, my target concept may be the concept *person*, identified as the concept expressed by the expression "person" in a given community C. If this is the starting point, then it counts as a priori that the concept *person* is the concept expressed by the expression "person" in C, for I simply stipulated that in order to identify the target concept. But I could now gather data about the way members of C use the expression "person", for example by presenting them with elaborate scenarios and asking them questions, or eliciting statements from them, involving the term "person". If I myself am a member of C, then my statements or answers to these questions have evidential force. It may have smaller or greater evidential force depending on the extent to which I am representative of C, and also depending on whether my answers are influenced by biases on my part. Thus, given these two important caveats, I can use my own judgements as empirical evidence regarding the concept expressed by "person".

---

[14] See footnote 5 above.

In principle, a target concept can also be identified without saying that it is the concept expressed by a given expression. For example, there may be a range of "personal attributes", i.e. properties we tend to attribute to persons. If we then stipulate that the concept *person* is that of a potential holder of personal attributes, then we could gather empirical evidence on how thinkers (in some community C) attribute this sort of property, and how they store this information. Thus, if friendliness, unfriendliness, laziness, smartness, honesty, viciousness etc were such personal attributes, we could look at how thinkers store such information in order to see how they identify the possessors of these properties over time. Again: the conceptual analyst's own behaviour has evidential weight to the extent to which they are representative of the target community, and to the extent to which they are not biased.

Thus, the classic armchair method to investigate concepts by considering hypothetical cases *can* be methodologically sound within the mentioned constraints. But the constraints are important and they give rise to legitimate concern about philosophers relying on their own judgements as evidence about a given concept. Such concerns have led to a wave of criticism of armchair methods in philosophy, criticisms that often do expose possible methodological flaws.

First, the conceptual analyst's own judgements may be biased because he or she has a stake in the outcome of their own investigation. Wishful thinking can impair one's judgement even in cases where one has no antecedent expectation or theoretical allegiance: for example through the wish to find something interesting. The situation is worse when philosophers rely on judgements concerning outlandish possibilities, that never occur in actual applications of the concept, so that it is unclear whether they are exercising an existing classificatory mechanism or rather an ability to extrapolate or creatively extend their conceptual repertoire.[15]

Secondly, conceptual analysts, in consulting their own judgements, may also fail to be representative of the target group. In a classic paper, Machery et al. question the evidential significance of theorists' judgements about proper names and their referents on the basis of cross-cultural variation in such judgements (Machery et al 2004). Obviously, when the analyst's own judgements, that are used as evidence, diverge from

---

[15] See Sytsma & Buckwalter 2016 and Machery 2017.

those of a significant part of the target group, then this undermines the evidential force of these data. Can an armchair conceptual analyst respond to this charge?

Concerns about representativeness require different responses depending on the exact target of investigation. Suppose the target is a concept employed by an individual, or perhaps by the individual and her local social group, and the conceptual analyst is either identical to the individual or a member of the local group. Then there will not be a problem with cross-cultural variation: the variation will simply show that other cultural groups do not have the same concept (or perhaps do not express the same concept by the expression in question) as the target individual or the target local group. (NB: even when there is no problem of representativeness, there may still be a problem of bias.)

The response will need to be different if the target of investigation is a concept of a certain group G and the variation in behaviour has been observed precisely across that G. [16] In this case, variation within G—be it cross-cultural or along different dimensions—is a problem, and calls for correction in the conceptual analyst's flawed procedure. Perhaps the identification of the target concept wrongly presupposes that group G is uniform in certain respect. Thus the analyst might correct the target by restricting the group and stick to the armchair method. Alternatively, she might stick to the target but gather data of a different sort, for example by conducting experiments on a properly representative sample of members of G.

On the view I have presented, the conceptual analyst's own armchair data and the data from controlled experiments on a representative sample, both constitute evidence of the same type: they measure the manifestations of those dispositions the analyst seeks to analyse. Armchair data, while less costly to produce, are less reliable, because they may be subject to biases or the researcher's failure to be representative of the target group. Proper experiments therefore produce superior data that allow more reliable inferences about the concepts employed by groups.

This does not mean, however, that we should ban armchair methods altogether. There could be cases where the superior evidence generated by controlled experiments,

---

[16] And this seems to be the case in the target of the studies to which Machery et al. intended to object in their early paper: the target here seems to have been a notion of reference assumed to be shared amongst all human societies. So the finding of cross-cultural variation was highly significant at least relative to that target of investigation.

as compared to armchair evidence, is not worth the extra cost of such experiments. There could be questions about concepts where it would be a waste of resources to conduct an experiment—simply because the evidence generated from the armchair is already good enough on those questions. For example, we may not need an experiment to check whether the concept *square* (as employed by such and such a group) entails the concept *equilateral*. Or whether it is part of possessing vague concepts (such as *heap*, *bald*, or *red*, as employed in a certain group) that undetectable differences do not make a difference. In some cases, the benefit of ruling out biases and ensuring the representativeness of the sample might be so slim as to be simply outweighed by the extra cost of conducting suitable experiments. Moreover, even before such an experiment is carried out, a researcher would be well advised to consider first the cheap and easy evidence that can be gotten from the armchair, before investing much more effort into an experiment. The armchair evidence such a potential experimenter already has before the experiment is not worthless.

Thus, to sum up my conclusions about empirical investigations regarding concepts we actually use: Depending on the exact question we are trying to answer, relying on the researcher's own judgements as evidence can be perfectly adequate. However, in case where there is a danger of researcher bias, or where there is variation across different thinkers in the target group, experimental methods are the obvious way to improve the evidential situation. There are both important continuities and important differences between experimental and armchair methods. Both are in principle sound empirical methods for investigating concepts within the limits pointed out above.

6. A priori Investigations of Concepts

Following the distinction from §3, a priori investigations about concepts rely on deduction and stipulation only. Of course, the status of deduction as a source of knowledge is an important area of research in itself. But for current purposes it is not unreasonable, I believe, to assume that deduction is a bona fide a priori source of knowledge. The same goes for stipulation: when I decide to investigate concepts that are governed by a certain principle, then it seems as uncontroversial as it is trivial that I know that these concepts are governed by that principle. Thus, any method that investigates the deductive consequences of stipulations seems in principle to be a sound

a priori method (of course, this does not rule out the possibility of deductive error or incoherent stipulations).

The challenge is rather to explain how mere stipulation and deduction could yield any philosophically interesting results. I believe that most of the philosophical interest of a priori results derives from the actuality or possibility of us using concepts with the features deductively examined. I shall use this section to distinguish four ways in which a priori results about concepts may be philosophically useful or interesting. The next section will provide illustrative examples of these ways.

First, there may be such a thing as a purely intellectual interest in making observations about the deductive consequences of some properties a concept may have. Thus, perhaps some philosophers are interested in the exact settings in which an equivalence principle concerning the concept of truth leads to the liar paradox, and their interest is completely independent of the possibility that a concept we actually use may be subject to that principle. I believe that that is a possibility, and that there is a lot of such non-trivial a priori knowledge that one could acquire. However, it is doubtful that this purely intellectual interest drives much philosophical research.

Secondly, we may have a desire to understand the concepts we have. A priori results about concepts with certain characteristics can derive their interest from a desire to understand ourselves and our own intellectual tools. For we may have (empirical) reason to believe that we actually employ concepts with characteristics more or less similar to those explored. (I suspect that even the most abstract logical investigations about concepts will typically derive in some way from such an assumption of empirical relevance.) Once we believe that we employ concepts with certain characteristics, we may want to explore the consequences of this.

Perhaps this is the time to come clean that exploring the deductive consequences of certain characteristics of concepts may include exploring these consequences under certain further empirical hypotheses. For example, suppose we are exploring the consequences of a tolerance principle governing vague concepts, which says that if the concept applies to an object, then it must also apply to any undetectably different objects. Then we will be interested in the consequences of the principle on the empirical assumption that a sorites series, in which neighbouring objects do not differ detectably, exists.

Thirdly and relatedly, then, our interest may be diagnostic. We may be interested in diagnosing problems that we suspect may arise from the concepts we employ. Again, given the empirical assumption that we use a concept with given characteristics, and perhaps given further empirical assumptions, we may be interested in working out what follows deductively. The results may show that a concept is inconsistent, that it is trivial, or perhaps that it is inconsistent with certain further assumptions we want to make.

Fourthly, our interest may be constructive. We may be interested in mapping out the properties even of concepts we do not currently use because we want to know what it would be like if we were to use them. The aim here may be to improve our conceptual repertoire by replacing or improving old concepts or adding new ones. Again, mapping out the consequences here may involve making further empirical assumptions.

To sum up, while there is plenty of non-trivial a priori knowledge about concepts to be had, there are three important ways in which this knowledge can become philosophically interesting in so far as these concepts are actually, or might potentially, be employed by thinkers.

7. Conceptual Analysis at Work

The debate about the problems of armchair methods in supporting empirical claims about concepts in my view distracts from some of the most important roles of conceptual analysis in philosophy. The a priori study of concepts independently of whether anyone ever uses these concepts, may be of limited interest. But by itself, the empirical study of the concepts actually employed by human thinkers looks more like a psychological or sociological, than a philosophical method.[17] In this section, I want to illustrate with concrete examples how all three methods of conceptual analysis: the two empirical methods (armchair and experimental) and the a priori method, can be harnessed to advance philosophical concerns, thereby uncovering the important roles conceptual analysis has in philosophy.

---

[17] I am aware that some philosophers do have doubts about the legitimacy of philosophy as a separate discipline precisely because they think philosophy is mainly concerned with empirical questions about concepts that would be much better addressed with the empirical methods of psychologists or sociologists. I hope this paper provides the germ of a historically more charitable conception of philosophy's subject matter.

In §6, I distinguished three types of philosophical aim for which a priori knowledge about concepts can be useful in so far as the concepts in question are, or might be, employed by us: (i) clarifying our own concepts, (ii) diagnosis of conceptual problems and (iii) conceptual improvement. Much of the most interesting philosophical work simply involves purely a priori reasoning about concepts that are of interest to us because they are concepts that we do use, are similar to concepts we use, are concepts that we might consider using, etc. In order to back up this claim, let me illustrate it with three examples.

The first example is that of research aimed at resolving the liar paradox. Researchers in this field are clearly engaged in elaborate *a priori* work concerning various precisely defined concepts (or languages), assessing and comparing their a priori derivable properties. These are indeed exercises in deduction from stipulative definitions, but they need not be as "easy" and as irrelevant as Jackson and Strawson suggest (see quotes above). Such research derives its interest from the fact that a concept we use, expressed by the English expression "true", seems to be subject to a problematic requirement, namely that it make true instances of the equivalence schema. Whether anyone employs a concept that is indeed subject to such requirements is an empirical question to be investigated by the kinds of empirical methods mentioned in §5. But many of those working on the liar think of themselves as engaged in the project of examining the a priori consequences of such requirements on a concept (e.g. Tarski 1956/1933). Many liar researchers operate on the assumption that the truth concept we have must be consistent (empirical assumption), so that the project of finding a truth concept that is not liar-susceptible ultimately contributes to the empirical project of specifying the exact correctness conditions of the truth concept we have (philosophical aims (i) and (ii) above).

Other liar-researchers may follow a different narrative: that of providing a stipulatively defined concept that shares certain features of the truth concept we use, so that the former could take the latter's role, but is an improvement in that it does not suffer from the liar problem. The latter group may further split into those who think of themselves as offering a Carnapian *explication* (cf. Carnap 1950), and those who think

they are offering a replacement (see e.g. Scharp 2007, cf. Eklund 2002).[18] All three types of liar researchers engage in philosophically interesting work, work that is largely stipulative-deductive conceptual analysis. The philosophical interest derives from the fact that the stipulatively defined and deductively examined concepts are in some respects similar to concepts we use, but better in other respects (aim (iii) above).

The example of research on the liar is just an example in which it is especially clear that the purely a priori aspect is not easy, not trivial, and not irrelevant. There are countless research questions in philosophy that involve, or can involve, examining candidate definitions and the properties that follow. For example, questions on vagueness, identity, personal identity, necessity, theories of truth (when the liar is not concerned), probability, persistence through time, etc.

My second example is research into moral concepts, such as the concepts *morally good*, *morally required*, *rational* or related moral concepts. Clearly, research in this area often involves mapping out the consequences of precise candidate definitions, or candidate defining principles, of such concepts. This is largely *a priori* work, involving stipulative definitions and mapping out deductive consequences, although the consequences being mapped may include deductive consequences that follow from candidate definitions or principles *together with* certain empirical assumptions.

The theoretical aims that moral theorists see themselves as pursuing, and for which these purely a priori considerations are meant to be useful, can vary. A stipulative definition of "good", "required" or "rational" and its purely *a priori* properties, may be interesting because "good", "required" or "rational" as we actually use it, means something very similar. The *a priori* properties of a moral concept with given application conditions may be interesting because the moral concepts we actually employ have these, or similar, application conditions (aim (i)).[19]

But the properties that follow *a priori* from certain stipulations about expressions or concepts may also be interesting if we are trying to decide which expressions or concepts it would be good to have. Many moral theorists regard moral theory as a more or less revisionary project. Mill's *Utilitarianism* (1861) can be read in this way (see

---

[18] Yet others, e.g. Tarski 1944, do not share the assumption that there is a unique concept expressed by the predicate "is true", but instead think that the expression may be ambiguous or may not have a definite meaning.

[19] This seems to be the outlook in Jackson 1998 and in Jackson & Pettit 1995.

Brink 2013, §10). Richard Brandt explicitly pursues the project of proposing "reforming definitions" (Brandt 1979, pp. 3–13). Peter Railton (1989) not only sees his own task as revisionary, he also claims that

> most philosophers in [the 20th] century … have wanted to avoid outright abandonment of value discourse and have shied away from claiming that our value judgements are systematically false, and so have opted for some sort of revisionism. (159).

These philosophers are not changing the topic or reaching results that are irrelevant to our actual moral terms or concepts, as Strawson and Jackson fear. Rather, they are considering and proposing refinements or improvements (aim (iii)).

The idea of reflective equilibrium, taken from Goodman (in the context of justifying deduction) and applied by Rawls (1971) in the area of moral and political theory, seems to point in a similar direction. Reflective equilibrium seems to be an ongoing, gradual process of conceptual development: theorists construct moral principles or definitions that predict as systematically as possible, our considered moral judgements. Then, as a result of reflection on such principles, we change our mind about some of the considered judgements. We revise the principles, reflect again on the considered judgements and change our minds again. Such a process is often discussed as a method of justification in moral epistemology. But it is also a process of gradual revision of our conceptual repertoire, one that moral theorists seem to want to set off and engage in. We are replacing our moral concepts by slightly different new concepts, re-evaluate and replace again.

Here too, working out a priori consequences of candidate definitions or principles, and perhaps working out their consequences in certain concrete or hypothetical situations, is highly relevant and philosophically interesting, *not* because the definitions or principles accurately capture our conceptual practice, but the interest lies precisely in the differences between the stipulated concepts we examine and the concepts we seem to have been using when arriving at our considered judgements (aim (iii)).

The third example is that of recent explicit proposals to "engineer" certain concepts in order to improve our conceptual repertoires and thereby effect social or political changes. A classic point of reference in the growing literature on conceptual engineering are Sally Haslanger's and others' proposals to "ameliorate" the concepts

of race or gender (Haslanger 2000, see also Cappelen 2017, ch. 1, for an overview). Again, certain candidate concepts are identified stipulatively, the a priori and a posteriori consequences of adopting such concepts are explored. For example, Haslanger proposes a certain stipulative definition of "woman" and considers (using empirical assumptions about human societies) what would be the social and political consequences of adopting such a definition. In general, conceptual engineering involves examining concepts, stipulatively identified, and making a case, on practical, moral or political grounds, for prescribing the use of some of the explored concepts, and to legislate that certain expressions already in use should be taken to express these concepts (aim (iii)).[20]

In my view, these three examples provide paradigm examples of the important joint role a priori and empirical conceptual analysis can play in philosophy.

7. Conclusion

In this paper, I have addressed an embarrassing dilemma for the method of conceptual analysis: either (i) it aims at empirical theses about the minds of individuals or groups, and is therefore highly unreliable when performed from the armchair, or (ii) it aims at a priori theses about concepts that are stipulatively defined and is therefore irrelevant to philosophical concerns.

My response to the dilemma has been to embrace both horns: after introducing a minimal account of concepts, I distinguished empirical from purely a priori questions about concepts. Within the empirical category, I distinguished quasi-empirical "armchair" methods of empirically examining concepts from experimental methods, but argued that both can be legitimate, each with its own advantages and limits. I also made a case for the importance of purely *a priori* conceptual analysis in philosophy. This defence of *a priori* conceptual analysis did not rely on especially ambitious or contentious epistemological assumptions: all it required was that we can generate *a priori* knowledge by making use of stipulation and deduction. The defence relied on a

---

[20] True, the method of exploration is not only purely a priori: in addition to purely deductive means, the investigation has to rely also on empirical information on what effects concept use has. But the relevant empirical information here is not information about which concepts the target expressions actually express, but rather on what would be the causal effects on individuals and groups if the target expressions were to express certain different concepts.

pluralistic picture of what philosophers are up to when they engage in conceptual analysis, ranging from the exploration of concepts that we have empirical reason to believe are like the concepts we use, to the considerations it would be good for us to have. [21]

References

Båve, Arvid 2019: "Concept Designation". *American Philosophical Quarterly* 56, pp. 331–44.

Bealer, George 1998: "Intuition and the Autonomy of Philosophy". In M. DePaul and W. Ramsey (eds.), *Rethinking Intuition*, Oxford: Rowman & Littlefield, pp. 201–39.

Bengson, John 2015: "The Intellectual Given". *Mind* 124, pp. 707–60.

Brandt, Richard 1979: *A Theory of the Good and the Right*. Oxford: Clarendon Press.

Brink, David O. 2013: *Mill's Progressive Principles*. Oxford: Clarendon Press.

Burge, Tyler 1979: "Individualism and the Mental". *Midwest Studies in Philosophy* 4, pp. 73–121.

Camp, Elisabeth 2015: "Logical Concepts and Associative Characterizations". In Margolis & Laurence (eds.) 2015.

Cappelen, Herman 2017: *Fixing Language*. Oxford: Oxford University Press.

Cappelen, Herman 2012: *Philosophy without Intuitions*. Oxford: Oxford University Press.

Carey, Susan 2009: *The Origin of Concepts*. Oxford: Oxford University Press.

Carnap, Rudolf 1950: "On Explication", pp. 1–18 in *Logical Foundations of Probability*, Chicago: University of Chicago Press.

Chalmers, David and Frank Jackson 2001: "Conceptual Analysis and Reductive Explanation". *Philosophical Review* 110, pp. 315–60.

Chalmers, David 2012: *Constructing the World*. Oxford: Oxford University Press.

Eklund, Matti 2002: "Inconsistent languages". *Philosophy and Phenomenological Research* 64, pp. 251–75.

Frege, Gottlob 1884: *Die Grundlagen der Arithmetik*. Breslau: Verlag Wilhelm Koebner.

Glock, Hans-Johann 2011: "A Cognitivist Approach to Concepts". *Grazer Philosophische Studien* 82, pp. 131–63.

Grice, H. Paul 1958: "Postwar Oxford Philosophy". In *Studies in the Way of Words*. Cambridge, Mass.: Harvard University Press 1989, pp. 171–80.

Grice, H. Paul 1987: "Conceptual Analysis and the Province of Philosophy". In *Studies in the Way of Words*. Cambridge, Mass.: Harvard University Press, pp. 181–5.

Haslanger, Sally 2000: "Gender and race: (What) are they? (What) do we want them to be?". *Noûs* 34, pp. 31–55.

Haug, Matthew C. (ed.) 2014: *Philosophical Methodology: the Armchair or the Laboratory?*, London: Routledge.

Horvath, Joachim 2018: "Philosophical Analysis: The Concept Grounding View". *Philosophy and Phenomenological Research* 97, pp. 724–50.

Jackson, Frank 1998: *From Metaphysics to Ethics: a Defence of Conceptual Analysis*. Oxford: Oxford University Press.

Jackson, Frank and Philipp Pettit 1995: "Moral Functionalism and Moral Motivation". *Philosophical Quarterly* 45, pp. 20–40.

Juhl, Cory and Eric Loomis 2009: *Analyticity*. London: Routledge.

Kauppinen, Antti 2013: "A Humean Theory of Moral Intuition". *Canadian Journal of Philosophy* 42, pp. 360–81.

Kornblith, Hilary 2002: *Knowledge and its Place in Nature*. Oxford: Clarendon Press.

Kripke, Saul 1975: "Outline of a theory of truth". *Journal of Philosophy* 72, pp. 690–716.

Laurence, Stephen and Eric Margolis 2003: "Concepts and Conceptual Analysis". *Philosophy and Phenomenological Research* 67, pp. 253–82.

Levin, Janet 2014: "Reclaiming the Armchair". In Chris Daly (ed), *The Palgrave Handbook of Philosophical Methods*. Basingstoke: Palgrave, pp. 448–77.

Margolis, Eric and Stephen Laurence 2007: "The Ontology of Concepts—Abstract Objects or Mental Representations?". *Noûs* 41: 561–93.

Margolis, Eric and Stephen Laurence 2014: "Concepts", *The Stanford Encyclopedia of Philosophy* (Spring 2014 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/spr2014/entries/concepts/>.

Margolis, Eric and Stephen Laurence (eds.) 2015: *The Conceptual Mind*. Cambridge, Mass.: MIT-Press.

Machery, Edouard, Ron Mallon, Stephen Nichols and Stephen Stich 2004: "Semantics, Cross-Cultural Style". *Cognition* 92, B1–B12.

Machery, Edouard 2009: *Doing without Concepts*. New York: Oxford University Press.

Machery, Edouard 2017: *Philosophy Within Its Proper Bounds*. Oxford: Oxford University Press

Matthews, Robert J. 1994: "The Measure of Mind". *Mind* 103, pp. 131–46.

Mill, John Stuart 1861/2003: *Utilitarianism*. In *Utilitarianism* and *On Liberty*, Mary Warnock (ed.), Oxford: Blackwell, pp. 181–235.

Millikan, Ruth 2000: *On Clear and Confused Ideas:An Essay on Substance Concepts*. Cambridge: Cambridge University Press.

Millikan 2011: "Loosing the Word-Concept Tie". *Proceedings of the Aristotelian Society Supplementary Volumes* 85, pp. 125–143.

Nolan, Dan 2009: "Platitudes and Metaphysics". David Braddon-Mitchell and Robert Nola (eds.), *Conceptual Analysis and Philosophical Naturalism*, Cambridge, Mass.: MIT-Press, pp. 267–300.

Papineau, David 2009: "The Poverty of Analysis". *Proceedings of the Aristotelian Society Supplementary Volume* 83, pp. 1–30.

Papineau, David 2014: "The Poverty of Conceptual Analysis". In Matthew C. Haug (ed.), *Philosophical Methodology: the Armchair or the Laboratory?*, London: Routledge, pp. 166–94.

Peacocke, Christopher 1992: *A Study of Concepts*. Cambridge, Mass.: MIT-Press.

Railton, Peter 1989: "Naturalism and Prescriptivity". *Social Philosophy & Policy* 7, pp. 151–74.

Rawls, John 1971: *A Theory of Justice*, Cambridge, MA: Harvard University Press. Revised edition, 1999.

Ryle, Gilbert 1949/2009: *The Concept of Mind*. Julia Tanney (ed.), London: Routledge.

Sainsbury, Mark and Michael Tye 2011: "An Originalist Theory of Concepts". *Aristotelian Society Supplementary Volume* 85, pp. 101–24.

Sainsbury, Mark and Michael Tye 2012: *Seven Puzzles of Thought and How To Solve Them*. Oxford: Oxford University Press.

Scharp, Kevin 2007: "Replacing truth". *Inquiry* 50, pp. 606–21.

Schroeter, Laura 2006: "Against *A Priori* Reductions". *Philosophical Quarterly* 56, pp. 562–86.

Soysal, Zeynep 2018: "Formal Analyticity". *Philosophical Studies* 175, 2791–811.

Stalnaker, Robert 2001: "Metaphysics Without Conceptual Analysis". *Philosophy and Phenomenological Research* 62, pp. 631–6.

Strawson, Peter F. 1963: "Carnap's Views on Conceptual Systems versus Natural Languages in Analytic Philosophy". In Paul Arthur Schilpp (ed.), *The Philosophy of Rudolf Carnap*. La Salle, IL: Open Court, pp. 503–18.

Strawson, Peter F. 1992: *Analysis and* Metaphysics. Oxford: Oxford University Press.

Sytsma, Justin and Wesley Buckwalter 2016: *A Companion to Experimental Philosophy*. Oxford: Blackwell.

Tarski, Alfred 1944: "The Semantic Conception of Truth and the Foundations of Semantics". *Philosophy and Phenomenological Research* 4, pp. 341–76.

Tarski, Alfred 1956/1933: "The concept of truth in formalized languages". In Alfred Tarski, *Logic, Semantics, Metamathematics*. Oxford: Clarendon Press, pp. 152–278.

Thomasson, Amie L. 2015: *Ontology Made Easy*. New York: Oxford University Press.

Williamson, Timothy 2003: "Blind Reasoning". *Proceedings of the Aristotelian Society Supplementary Volume* 77, pp. 249–93.

Williamson, Timothy 2007: *The Philosophy of Philosophy*. Oxford: Blackwell.

Williamson, Timothy 2018: *Doing Philosophy*. Oxford: Oxford University Press.

Yablo, Stephen 1993: "Is Conceivability a Guide to Possibility?". *Philosophy and Phenomenological Research* 53, pp. 1-42.

Zalta, Edward 2001: "Fregean Senses, Modes of Presentation, and Concepts". *Philosophical Perspectives* 15: pp. 335–59.