

Why the Canberra plan won't help you do serious metaphysics

Raamy Majeed¹ 

Received: 5 November 2016 / Accepted: 5 May 2017 / Published online: 23 May 2017
© The Author(s) 2017. This article is an open access publication

Abstract Jackson (From metaphysics to ethics, Oxford, Oxford University Press, 1998) argues that conceptual analysis plays a modest, albeit crucial, role in ‘serious metaphysics’: roughly, the project of demystifying phenomena we take to be mysterious by locating them in the natural world. This defence of conceptual analysis is associated with ‘the Canberra Plan’, a philosophical methodology that has its roots in the works of both Lewis (J Philos 67(13):427–446, 1970, Australas J Philos 50:249–258, 1972) and Jackson (Monist 77:93–110, 1994, 1998). There is, however, a distinction to be drawn between conceptual analysis, as it is typically employed in the Canberra plan, and a version of it defended by Jackson himself. In this paper, I elucidate this distinction, and employ examples from the history of science to argue the use of the former, but not the latter, incurs certain problems of conceptual change. Moreover, I also argue neither can be used to undertake serious metaphysics—the former because of the aforementioned problems, and the latter due to the machinery it employs to solve them.

Keywords Canberra plan · Conceptual analysis · Conditional analysis · Conceptual change · Best-deserver theories · Meta-metaphysics

1 Serious metaphysics

Metaphysics is usually regarded as being the grand old project of discovering what there is and what it is like. In his seminal book, *From Metaphysics to Ethics: A Defence of Conceptual Analysis*, Jackson (1998) begins, however, by noting that metaphysi-

✉ Raamy Majeed
mra37@cam.ac.uk

¹ Faculty of Philosophy, University of Cambridge, Cambridge, United Kingdom

cians actually concern themselves not just with taking an inventory of the world, but rather providing a comprehensive account of some subject-matter in a limited amount of more basic notions. ‘Serious metaphysics’, then, is understood in terms of what he dubs *the location problem*: showing how to locate the mystifying subject-matter in the natural world. To elaborate, when confronted with facts or properties at a mysterious level, what we do as metaphysicians is attempt to find a place for them in a world we like: facts or properties at a non-mysterious level. For example, if we are physicalists, we attempt to ‘locate’ the mysterious matter, like the mind, freewill, morality, colours, etc., in a physical reality.

Jackson’s interlocutors don’t quibble with the claim that serious metaphysics involves attempting to solve the location problem. Nor the claim that this problem can only be solved via *the entry by entailment thesis*, i.e. by showing that the non-mysterious facts entail the mysterious facts; thereby illustrating how they are ‘located’ in the non-mysterious facts. The disagreement rests on whether we require conceptual analysis as a tool for establishing this thesis. Jackson argues that it plays a modest, albeit crucial, role whereas his detractors argue otherwise. Typically, they cite Quine, Putnam, Kripke—attacks on apriority-cum-analyticity, causal theories of reference, a *posteriori* and metaphysical necessities—as a way of demonstrating that we can support the thesis *sans* conceptual analysis.¹ For his part, Jackson argues that these, by and large, are red herrings, and we require conceptual analysis to identify our subject-matter. That is, we need to identify ‘our ordinary conceptions’ of the mysterious facts before we can begin any investigation as to whether they are entailed by the non-mysterious ones.

My own sympathies lie here with Jackson. This paper, subsequently, concerns itself not with any challenges to conceptual analysis along the what-are-now traditional lines. Instead, it makes the relevant concessions to Jackson, and aims to address certain problems concerning conceptual change, which I take to be manifest in some contemporary versions of conceptual analysis associated with Jackson. The focus, in particular, will be on ‘Canberra plan’ style analyses, which have their roots in both Lewis (1970, 1972) and Jackson (1994, 1998), as well as another version of conceptual analysis defended by Jackson himself.

The structure of the paper is as follows. In the first part of the paper, I shall provide a brief exposition of the Canberra plan (Sect. 2), and demonstrate, using examples from the history of science, how analyses of this type face certain problems of conceptual change (Sect. 3). I argue that one of these problems is doubly troublesome because it also inhibits our ability to reliably perform serious metaphysics. In the second part of the paper, I explain how the version of conceptual analysis actually spelt out by Jackson is immune from the problems of conceptual change (Sect. 4), but only at the expense of not being able to settle questions concerning which mysterious facts are entailed by the not-so mysterious ones (Sect. 5). I take the conjunction of these two parts to show that neither version of conceptual analysis discussed can prove helpful when it comes to settling serious metaphysical debates.

¹ For an introduction, see the objections by Block and Stalnaker (1999), and the response from Chalmers and Jackson (2001).

A caveat. Most Canberra planners, as well as Jackson, are deflationists about concepts in the sense that they take talk of concepts and conceptual analysis to be that of words (especially concept-terms), and the possible situations covered by their use. As Jackson notes, the talk is maintained only in deference to the traditional terminology. This paper will follow suit with the additional proviso that the notion of conceptual change will be understood as the employment of the same concept-terms in referring to distinct subject-matter. That being said, most of the issues discussed will, I take it, apply even if one assumes an inflationist framework. I shall, however, leave it up to inflationists themselves to work out the finer details of how these issues relate to their particular theories.

2 The Canberra plan

‘The Canberra plan’ is an umbrella term for various types of analysis that stem from Lewis and Jackson, and was practiced by philosophers in, and associated with, the Australian National University—located, of course, in Canberra. These analyses differ both superficially and substantially with each other, as one might expect, but they are said to share the same general strategy for settling metaphysical debates. Nolan (2009) calls this the ‘Canberra plan two-step’, and formulates it as follows:

Step 1: Give a functional analysis of the theoretical terms. E.g. *x* is whatever occupies the *x*-role.

Step 2: Look to our theory of the world to see what occupies the *x*-role.

Not all applications of the Canberra plan are intended to solve the location problem, though some are.² Furthermore, it is easy to see how these steps can provide a solution by lending support to the entailment thesis. Provided the non-mysterious basic facts are thus-and-so, we can show that these facts entail the mysterious facts. This is best seen by example.

Take the mysterious facts to be those concerning pain. We can give a Canberra plan style analysis of ‘pain’ in the following way:

- (1) Pain = whatever occupies the pain role. (By analysis)
- (2) C-fibre firing = occupant of the pain role. (By science)
- (3) ∴ Facts about C-fibre firing entail facts about pain. (Entailment thesis)

This borrows from the physicalist analysis of mental states offered by Lewis (1972, 1994). The difference concerns (3). Lewis was interested in offering an identity theory, viz. that mental states, including pain states, are identical to physical states. The demands on us with respect to the location problem are less stringent, as we need not establish an identity claim. For our purposes, it will suffice to note that the conjunction of steps (1) and (2) show how certain non-mysterious facts entail the mysterious facts concerning pain.

² Several applications of the Canberra plan, including by Lewis and Jackson, don’t aim to demystify mysterious phenomena. Rather, they focus on phenomena we may expect to find a place for in a naturalistic picture of what our world is like, e.g. solidity, motion, and belief, but which we refer to by terms that don’t (yet) belong to the languages of the natural sciences. I am indebted to a referee for this journal for this point.

The point of interest for us here is step (1), as it is that is germane for conceptual analysis. Canberra planners typically explain this step by applying the Ramsey–Carnap–Lewis (RCL) method for defining theoretical terms.³ The general idea is that we take a postulate of a scientific theory, call it T , and we define its theoretical terms, the t -terms, in terms of the theoretical roles the entities they refer to are supposed to play within that theory.

Here the theoretical roles are typically represented by a ‘Ramsey sentence’: an existentially quantified sentence, which replaces all the t -terms with variables, but keeps the o -terms: the non-theoretical original old terms of that theory. So for a postulate $T(t_1 \dots t_n, o_1 \dots o_n)$, we define its Ramsey sentence as ‘ $\exists x_1 \dots \exists x_n T(x_1 \dots x_n, o_1 \dots o_n)$ ’. Formally, this tells us that there is at least one actual realization of T . But implicitly, it also defines the t -terms in terms of the variables and o -terms.⁴ It is this feature, which is exploited by the Canberra planners.

According to Lewis (1970), what this yields is a functional analysis of the theoretical terms, as the Ramsey sentence defines them in terms of the functional roles their referents occupy within that theory. Jackson (1998) makes two crucial additions. First, he extends this approach beyond scientific theories to any possible theory, including folk theories. Second, and consequently, he takes the theoretical roles to be whatever the theory deems relevant—even if they surpass the causal roles typically associated with scientific theories. Jackson, for his part, still takes the theoretical roles to be functional, but he no longer supposes that they have to be causal. For example, a folk theory of pain might declare the relevant theoretic roles to be non-causal descriptive roles, like *revelation*: having certain intrinsic-cum-essential features of experiences, in this case pain experiences, revealed to you experientially.⁵

All this concerns conceptual analysis in the following sense. Lewis, Jackson, and the Canberra planners all agree that the central goal of conceptual analysis is precisely to make these roles explicit. Recall, according to Jackson, the broad aim of conceptual analysis is to identify our subject-matter. This, as we now see, involves identifying the relevant theoretical roles.

To elaborate, when it comes to the subject-matter of the sciences, presumably, the causal roles the entities posited are supposed to play can be read off the scientific theories themselves. Folk theories, however, aren’t written down. Therefore, it is our job, as philosophers, to make explicit what these theoretical roles are. Only then can we take our cues from the RCL method, and define their subject-matter via these roles.

It bears noting that making folk roles explicit might be construed as a far cry from the standard conception of conceptual analysis, which sees us attempting to define necessary and sufficient conditions for something to fall under a given concept. Nevertheless, as Jackson notes, our endeavour still bears a resemblance to the old one in that our work is still a priori and done from the armchair. According to the Canberra

³ See Ramsey (1931), Carnap (1963) and Lewis (1970, 1972).

⁴ For detail see Lewis (1970). Also see Jackson (1998) and Nolan (2009) for a summary.

⁵ The term is due to Johnston (1992). See Lewis (1995) for a discussion of why revelation frustrates physicalist-cum-functional attempts to reduce qualia.

planners, we collect all the platitudes we can about our folk concept, and we then add them together to form the relevant theoretical roles.⁶

There is also another similarity; one that gets rarely mentioned, perhaps, owing to its obviousness. That is, like the old conceptual analysis, the aim here, ultimately, is finding reference-fixing devices—be they conditions, roles etc.—which determine the applicability of our terms. On standard conceptions of conceptual analysis, a concept-term applies to something when that thing satisfies certain necessary and sufficient conditions. When it doesn't, the term simply fails to refer. Likewise, according to the Canberra plan, a concept-term refers to something when that thing occupies the relevant roles. When the relevant roles aren't occupied, our term fails to refer. Here eliminativism is seen to be the better option, not location.

3 Problems of conceptual change

The last point helps spell out the problems of conceptual change. According to the RCL method as defined so far, the theoretical terms only refer if the theory is a true theory, i.e. if its theoretical roles are actually occupied.⁷ Otherwise, these terms will fail to refer. Subsequently, when we abandon a theory, we should abandon its theoretical terms as well. If we don't, we have simply changed the concept and disguised this by using the same term.

To elaborate, for a theoretical term '*t*', which is part of its term-introducing theory *T*, the roles *t* is said to occupy according to *T* fixes the reference of '*t*'. Accordingly, if we abandon *T*, we should then also forgo the claim that '*t*' has its reference fixed. We should, in effect, be eliminativists about *t*. Now suppose we replace *T* with a newer theory *T**, but we still continue to use '*t*'. If we do so, the reference of '*t*' will be fixed by the roles its referent is supposed to play in *T**. Subsequently, '*t*' would appear to have changed its reference when we switch from *T* to *T**. Our term '*t*' no longer refers to *t*, it refers to some distinct set of entities *t**.

This in itself is fine. We are allowed to use *t*-terms as we see fit. Moreover, on pains of coming up with new terms every time a theory has been discredited, we may very well opt to keep the *t*-terms of abandoned theories. But if we do choose to keep '*t*' when we switch from *T* to *T**, we should do so whilst recognizing that we are using '*t*' to refer to a distinct set of entities. That is, in the original case, we use '*t*' to refer to the entities that occupy the roles specified in *T*, and in the second case, the entities that occupy the roles in *T**. Insofar as '*t*' is a concept-term, then, switching from *T* to *T** while retaining '*t*' would be tantamount to conceptual change.

The problem with this implication of the RCL method is that it flies in the face of our actual terminological use. We often keep the theoretical terms of abandoned theories. What's more, we don't take this to mean we have changed the concept. This is

⁶ For detail on these platitudes, see Nolan (2009) and Donaldson (2015).

⁷ Lewis's modification of the Ramsey–Carnap method actually adds a qualification that no longer requires the theory to be a true theory. As I explain shortly, he does so precisely to pre-empt problems like the problems of conceptual change.

evident from the use of scientific terms themselves. For example, the atomic theory of Dalton has been widely discredited on grounds that it took atoms to be both indivisible and indestructible. Insofar as we employ the terminology of the RCL method, we can say that Dalton's theory gets the theoretical roles wrong. Now, if we follow the RCL method—as described so far—we should, ipso facto, abandon the atomic theory and be eliminativists about atoms. This, however, goes against the data provided by the history of science. The preceding atomic theories, e.g. of Thompson, Rutherford, Bohr etc., kept the theoretical term 'atom'. Moreover, scientists don't typically construe these theories to be about some distinct set of entities. These theories, as well as the ones we have around today are regarded as being about the same stuff, namely atoms. Explaining this data is a problem for advocates of the RCL method, as the method, presently understood, entails that scientists have simply changed the concept and are now referring to something else altogether.

Note, technically speaking, there are in fact two related problems here. The first concerns how our terms get to refer when we abandon the term-introducing theories with which we define our terms. In other words, how do we avoid eliminativism each time a term-introducing theory is shown to be false? This problem doesn't concern conceptual change, as such. The second problem is about how two (or more) distinct theories, which share the same t-terms, but define them differently, can be said to be referring to the same set of entities. This problem is better described as being about conceptual change.

How we refer to these problems doesn't matter, but I prefer to discuss them, more or less, together for the following reasons. While the two problems are distinct, the second piggybacks on the first. That is, we only get worries about conceptual change in the second instance because it is an implication of the RCL method that we get eliminativism if the term-introducing theory is false. Moreover, and not surprisingly, the proposed solution to the first offered by the Canberra planners also proves to be a way addressing the second problem.

The standard solution to both problems, of course, is to invoke best-deservers. The idea, roughly, is that if nothing occupies all the theoretical roles of the original term-introducing theory, we identify the referents with 'best', 'imperfect' or 'near-enough' deservers: roughly, whatever occupies enough of the roles. What is regarded as the prime example of this is Lewis's (1995) analysis of 'qualia'. Here Lewis grants the dualist that the folk roles for phenomenal concepts involve a role that cannot be accommodated by materialists, i.e. revelation. Nonetheless, he argues that materialists can still be realists about qualia because they can accommodate the other roles associated with the folk concept, i.e. causal-cum-functional roles, such as having certain discriminatory powers.

The best deserver-approach tends to be employed without much by way of elucidation. I suspect this is because the necessary work was done earlier, in Lewis's (1970) original modification of the Ramsey–Carnap method for defining theoretical terms. Lewis warns that theoretical terms ought not to be defined by the whole theory precisely because if any part of the theory is false, this will mean that our t-terms will fail to refer. Instead, any given t-term is to be defined as the unique entity, if any,

which satisfies *most* of the theoretical roles.⁸ As to the question of how many count as enough, Lewis concedes that “There is no settled answer to the question ‘how good is good enough?’” (1995: p. 140).

This strategy is also exploited by Jackson (1998) to explain the distinction between entities we want to locate and ones that we want to eliminate. We are eliminativists about phlogiston because phlogiston theory gets most of the theoretical roles wrong, whereas we are realists about atoms because original atomic theories got enough of the relevant roles right. Since there are near-enough deservers of the term ‘atom’, the term gets to refer. Hence, we get to keep the term, and we do so, crucially, without changing the subject.

Since Lewis’s (1970) modification of the Ramsey-Carnap method incorporates the best-deserver strategy, the Canberra planners adoption of the RCL method, subsequently, should be understood in light of it. A term according to the Canberra plan, then, refers only if the relevant theoretical roles of the term-introducing theory are occupied—where *the relevant roles* need not be understood as all the theoretical roles; they can be enough of them.

This helps address the problems concerning conceptual change in the following sense. We can account for the continuity of reference when we switch from theories T to T^* by noting that both theories have enough roles in common. More specifically, while we abandon T , we suppose that ‘ t ’ refers because T got most of the theoretical roles correct. So we don’t end up being eliminativists about t . Now, insofar as T^* has enough theoretical roles in common with T , we not only get to be realists about t , but we get to claim that the entities referred to by ‘ t ’ in T are the same as those referred to by that term in T^* . Thus switching from T to T^* whilst retaining the use of ‘ t ’ doesn’t involve the use of the same term to refer to a distinct set of entities, which are defined in a new way. In other words, retaining the same concept-term across theories doesn’t entail concept change.

On this picture, charges of conceptual change can only be laid if replacement theories don’t share enough theoretical roles with the theories they replace. In which case, the (modified) RCL method suggests, *prima facie* correctly, that any persistent use of the terms of old abandoned theories would be tantamount to conceptual change. Therefore, either way, the method seems to deliver the correct results.

The modified RCL, I grant, does save Canberra planners from charges of conceptual change in some instances, e.g. in their analysis of ‘qualia’. What I contest is whether the best-deserver approach can be invoked to explain all the historical data, which are found to be troubling. Here it is instructive to carefully examine some actual examples from the history of science for the following reason. The best-deserver approach is, *prima facie*, very plausible, and has proved to be an effective way to ward off sceptical responses to the Canberra plan which concern conceptual change. Subsequently, the real force of any challenge to it, I take it, comes not from whether it works as a problem-

⁸ Lewis (1970) takes all theoretical roles to be equally weighted, which is presumably why ‘imperfect-deservers’ are whatever occupies most of these roles. Papineau (1996: p. 16), by contrast, suggests that a better strategy would be to weight the theoretical roles with regard to their ‘definitional importance’, and to take the referent to be whatever, if any, which satisfies the most important ones. Problems with this strategy are discussed by Stich (1996) and Braddon-Mitchell (2005: §2.2), amongst others.

fixer in abstract philosophical debates about defining theoretical terms, but whether it adequately captures what the present debate is really supposed to be about: cogent real-life cases where we seem to be referring to the same subject-matter despite there being significant differences between the term-introducing theories and those which we presently adopt. Let us begin, then, by examining one of Jackson's own examples, viz. atoms.

A careful examination of the etymology of the term 'atom' gives us enough evidence to doubt whether replacement atomic theories retain enough of the theoretical roles of those they replace. According to the standard narrative, twentieth Century atomic theories, from that of Dalton onwards, are a continuation of those of seventeenth Century mechanistic philosophers, e.g. Descartes, Gassendi and Boyle, who in turn are seen to revive ancient Greek atomism. Historians of science are now challenging this picture. Alan Chalmers (2009), for example, argues that the 'philosopher's atom', viz. stone-like bits of matter that have no structure, are unchangeable, experientially unverifiable, collide in the void to form macroscopic objects, and which are the ultimate constituents of the world, are vastly different to the 'scientist's atom': structures, which are changing, susceptible to empirical verification, and no longer the ultimate constituents of reality.⁹

Perhaps the best-deserver response to conceptual change is more plausible if we restrict our domain to twentieth Century atomic theories. But even here, the issue is much more controversial than Canberra planners typically assume. Dalton's atomic theory, as mentioned, gets two key theoretical roles wrong. But did it have enough theoretical roles in common with the atomic theories that followed? Arguably not. Thompson's 'Plum-Pudding' model, for instance, gave Dalton's basic, albeit weighted, atoms a structure: a positively charged body with lumps of negatively charged electrons in its body. Rutherford and Bohr replaced this body with empty space, and posited something that looked like a solar system where planets orbit a sun, i.e. a set of negatively charged electrons orbiting a nucleus comprising of positively charged protons and zero-charged neutrons. Further still, the quark model of Gell-Mann and Zweig saw the electrons and neutrons, amongst other hadrons, being formed by more elementary particles called quarks, and so on.

This outline, of course, is much too brief to warrant a definite negative answer to our question. My view is that it does motivate scepticism about each given twentieth Century atomic theory retaining enough of the theoretical roles of its predecessor. However, this is controversial, for it is not hard to come up with examples in the form of roles that are shared amongst these atomic theories, e.g. that atoms are very small bits of matter that can react with each other to form certain chemical compounds, which in turn can decompose through other known chemical reactions, and so on. Adjudicating the issue proves difficult, and not least because how many roles count as enough, by Lewis's own admission, is vague. The central difficulty, however, is that the best-deserver approach is one that is, though routinely employed, severely under-developed. Lewis, for instance, says that it is "hard to analyze, but easy to understand" (1972: p. 252). Subsequent clarifications have only concerned how to select amongst

⁹ This line of argument isn't entirely new, e.g. see [van Melsen \(1952\)](#) and [Hey \(1934\)](#) before him.

a set of competing near-enough deservers, and even here, as Nolan (2005) notes, a systematic strategy remains absent. Luckily, the RCL method itself, which has been fleshed out more thoroughly, can be brought to bear on our question.

As noted earlier, the theoretical posits of a given theory, according to the RCL method, are defined functionally. This is significant because the theoretical roles are understood as functional roles whose (joint) occupation is *constitutive* of any given theoretical posit. For example, on the functional analysis of ‘pain’, pain *just is* whatever occupies the relevant functional roles for pain. Likewise, for any theoretical posit x , x *just is* whatever occupies the relevant theoretical-cum-functional roles for x within the x -positing theory. The best-deserver approach inherits this functional analysis. The only difference is that the approach relaxes the criterion that all the theoretical roles need to be (jointly) occupied for something to count as the relevant theoretical posit. Consequently, though Lewis modifies the RCL method by invoking best-deservers, the relevant roles are still those whose occupations are constitutive of the theoretical posits.

This helps address at least some potential counterexamples. Whether a given atomic theory refers to the same theoretical posits as its predecessor depends not on whether we can, in hindsight, point to features they have in common, but on whether they share enough theoretical roles. For Lewis, what counts as the relevant theoretical roles are those that feature in the Ramsey sentence of the term-introducing theory. This strikes me as unnecessarily restrictive, as the defender of the Canberra plan would have to point to not just examples of enough shared roles across the Ramsey sentences of various atomic theories, but to those going as far back as that of the term-introducing theory; say Dalton’s, or worse, the original atomic theory of Democritus. Suppose, instead, that the comparison class is the theoretical roles specified in the Ramsey sentences of any given two atomic theories. We can then say that these two theories are about the same subject provided they have enough theoretical roles in common, and not if they don’t. Whilst not entirely uncontroversial, what the aforementioned etymology of ‘atom’ demonstrates, arguably, is that for some of the atomic theories we wish to compare, even within the twentieth Century, these will fail to have enough theoretical roles in common. This should be worrisome to Canberra planners who want to avoid conceptual change.

There are other terms whose historical uses don’t fit with the modified RCL method. Some come from psychology. Changes in the Diagnostic Statistical Manual of Mental Disorders (DSM) by the American Psychiatric Association are prime candidates for the best-deserver approach in action. A quick comparison between DSM-IV and DSM-5, for instance, reveals an array of slight changes in pre-existing classifications, which are, *prima facie* at least, about the same subject-matter. Nonetheless, the changes from DSM-I&II to DSM-III couldn’t be any more different.

Psychologists in the sixties and seventies were, by and large, shying away from psychodynamics (psychoanalysis) and heading towards behavioural therapy. The success of the latter, coupled with new discoveries about the brain, lead not only to significant changes in the treatment of disorders, but also in how these disorders were themselves understood. This precipitated an extreme shift in the classification of disorders in DSM-III, published in 1980, from disorders defined via psycho-dynamic roles with a

focus on etiology, to those defined by behavioural-cum-physiological roles focussing on symptoms.¹⁰

One example is the classification of autism from a form of childhood schizophrenia, which was in turn understood as a form of psychosis caused by ‘refrigerator mothering’, to a disorder in its own right that was to be diagnosed via behavioural roles, e.g. lack of responsiveness to other people, resistance to change, and ‘peculiar’ interests in inanimate objects. This seems to be a case where the theoretical roles, whose occupation was assumed to be constitutive of autism, are fundamentally distinct to those of prior theories. Hence, the newer autism theory of DSM-III can’t be said to be talking about the same subject-matter as its predecessors according to the modified RCL method. Yet psychologists in the eighties presumably didn’t think they had begun to treat something besides autism.¹¹

Potential counterexamples loom here too, and it is instructive to briefly address them by way of further clarifying the problem at issue. For instance, one may expect that conceptions of autism in any of the DSMs, on account of them being tools for diagnosis, must have enough diagnostic criteria in common, e.g. they must specify enough shared symptoms such as a lack of responsiveness to other people, and a gross impairment in communication skills. One difficulty with motivating this counterexample is that explicit diagnostic criteria, as noted in the introduction to DSM-III, are often lacking in DSM-I and DSM-II. Whilst the aforementioned symptoms are explicitly stated in DSM-III, autism only gets two mentions in DSM-II, one to ‘autistic behaviour’, the other to ‘autistic thinking’. Nonetheless, the Canberra planner could insist that these behavioural and cognitive traits are understood precisely as the symptoms stated in DSM-III, and thus the two manuals have enough (implicit) diagnostic criteria in common.

That DSM-II and DSM-III presume some of the same symptoms in their characterizations of autism is surely plausible, and I do not wish to deny it. Nor do I wish to deny that psychologists, in practice, employ such shared symptoms in order to assist with their diagnoses. But symptoms aren’t the same as diagnostic criteria, as two diseases, mental or otherwise, can share the same symptoms whilst being diagnosed differently. Such criteria are heavily theory-laden. This is well documented, and exemplified by the case at hand.¹² On account of the psychodynamic theory, which informed DSM-II, the aforementioned behavioural symptoms, and those like it, weren’t thought of as being constitutive of autism, and ergo not criteria that could offer a definitive diagnosis, whereas they were in DSM-III given the push towards behaviourism. Subsequently, the theoretical roles, which would feature in the Ramsey sentences of autism in DSM-

¹⁰ These changes are well documented, e.g. see [First \(2012\)](#). Also see [Mayes and Horwitz \(2005\)](#) for an alternative explanation for the changes vis-à-vis pressures from health-insurance companies and increasing government influence.

¹¹ Coincidentally, the classification of autism is one of the central sources of controversy surrounding the DSM-5, for it eliminates the four subcategories of autism, including Asperger’s disorder, in favour of one disorder rated on a scale of severity. One could, subsequently, argue that we get fundamental distinctions in the theoretical roles here as well. For a discussion, see [Baker \(2013\)](#).

¹² It is also exemplified with regard to other mental disorders, e.g. see [Hacking’s \(1995\)](#) discussion of multiple personality disorder. A general exposition of how theory shaped diagnostic criteria in the shift from DSM-II to DSM-III can be found in [First \(ibid\)](#).

II and DSM-III, won't have enough in common for them to be about the same disorder according to the best-deserver approach.

It appears, then, that there are quite a few distinct linguistic communities that employ the same concept-terms, and on the face of it to refer to the same entities, whilst recognizing significant changes in the theoretical roles they associate with these concepts. The modified RCL, and by extension the Canberra plan, lacks the resources to explain this.

That's not all. The first of the problems concerning conceptual change is doubly troublesome for Canberra planners, as it entails that they can't reliably perform serious metaphysics. As we have seen, if nothing occupies the relevant theoretical roles, we have to declare the theoretical posits to be non-existent. So rather than 'locating' the entities in nature, we eliminate them instead. This in itself needn't be problematic, but coupled with the kinds of things term-introducing theories are wrong about gives us reason to worry.

To illustrate this, assume the aforementioned etymologies of 'atom' and 'autism' are both accurate. That is, let's grant the term-introducing theories for both concept-terms got most of their theoretical roles wrong, i.e. nothing occupies them. It would then turn out that these terms no longer refer. The upshot of this is that we would, in effect, have to be eliminativists about both atoms and autism.

Now, a strategy for performing serious metaphysics, *qua* solving the location problem, must allow us to eliminate certain mysterious phenomena, e.g. ghosts, ectoplasm, and phlogiston. But an upshot of employing such a strategy shouldn't be that we end up eliminating the kinds of things we have good independent grounds to posit, and thereby wish our analyses to retain. Presumably, atoms, certain psychological disorders etc., are precisely the kinds of things we do wish to retain. Thus the problem is that even though we are still, technically speaking, doing serious metaphysics, we aren't doing so reliably, as we end up with an ontology that is too sparse.

It is worth noting that Canberra planners often end up with sparse ontologies, e.g. Lewis's 'Australian materialism', which contends that qualia *just are* whatever occupies the relevant causal-cum-functional roles, and Jackson's moral functionalism, which locates ethical properties amongst the natural. The aforementioned problem is that we end up with an ontology that is radically sparse; much sparser than any of the typical Quinean desert landscapes. We lose not just possible objects and other kinds of mysterious entities, but also the scientific posits in which the naturalistically minded Canberra planners wish to locate the mysterious. The road to Canberra may involve desert landscapes, but it certainly isn't barren. A metaphysics beset by conceptual change, by contrast, is.

4 Jackson's conceptual analysis

Jackson's account of conceptual analysis arguably avoids the aforementioned problem of conceptual change. To be clear, the specific aim of conceptual analysis, as stressed by Jackson, is the same as that of the Canberra plan: we aim to make explicit the theoretical roles implicit in our folk theories. The strategy he employs, however, is distinct. Here we discover the theoretical roles by partitioning possible cases into ones

in which facts about the relevant (mysterious) entities obtain and ones in which they do not.¹³

Lets say we want to analyze the folk concept of ‘qualia’. The strategy is to think of various possible cases where facts about qualia obtain, and ones where they don’t. In a sense, what we are thereby doing is partitioning possibilities into ones where the term ‘qualia’ applies, and ones in which it doesn’t. Suppose, for example, the actual world is as such that we discover that there are no nonphysical ‘spooky’ states, e.g. states of which revelation holds true. Would we still claim that we have phenomenal experiences, like pain experiences? How we answer these kinds of questions help make explicit the theoretical roles we associate with ‘qualia’.

If we answer yes, we see the theoretical role alluded to by dualists, e.g. revelation, isn’t one of the theoretical roles of the folk concept. (Or at least, it isn’t a theoretical role the folk theory takes to be essential). If we, however, answer no, we see that such dualist roles are part of the folk concept. In this way, we get to discover the relevant theoretical roles of our folk theory of ‘qualia’ by the partitioning of possible cases. According to Jackson, this type of partitioning of possible cases is constitutive of conceptual analysis.

This way of analyzing ‘qualia’ is the motivation behind the so-called conditional analysis of phenomenal concepts. The similarity is especially apparent in Braddon-Mitchell’s (2003) analysis, which can be formulated as follows:

- (CA1) If there are nonphysical states of the relevant type in the actual world, ‘qualia’ refers (rigidly) to these states.
 (CA2) If there are no nonphysical states of the relevant type, but there are physical states of the relevant type, ‘qualia’ refers (rigidly) to these physical states.

Here what counts as nonphysical and physical states of the relevant type is an in-house debate to be settled within dualism and physicalism. For present purposes, we can note that the former are those entities, which occupy the descriptive roles the best dualist theory associates with the folk concept, whereas the latter are those that occupy the roles associated by the best physicalist theory. As we have seen, the dualist roles may include revelation, whereas the physicalist ones might be causal-cum-functional roles—where these roles are determined either empirically or by common-sense.

What this analysis does, in effect, is exploit the partitioning of possible cases in order to accommodate the distinct sets of roles, which may be relevant for determining the reference of ‘qualia’. We see that in possibilities, viz. worlds considered as actual, where the dualist roles are satisfied, ‘qualia’ refers to these (nonphysical) occupants.¹⁴ Whereas, in possibilities where the dualist roles aren’t occupied, but the causal-cum-functional ones are, ‘qualia’ refers to these (physical) occupants. The immediate upshot is that neither the standard dualist nor physicalist analysis of ‘qualia’ is technically

¹³ Note: some Canberra planners do employ this strategy. Jackson (1998) himself appears to switch between the two, e.g. chapter (2) vs. chapter (5). Nonetheless, the Canberra plan two-step, as typically understood, leaves out the strategy of partitioning possible cases in favour of collecting platitudes.

¹⁴ Possibilities can be carved up based on worlds considered as actual or counterfactual, which is also where the rigidity/non-rigidity of a term becomes salient. These issues are orthogonal to those covered in this paper. See Majeed (2014) for a detailed exposition of how they pertain to the conditional analysis of ‘qualia’.

correct. Rather, the correct analysis of ‘qualia’ is captured by the conjunction of (CA1) and (CA2).¹⁵

We can either take this to mean that the concept has two theoretical roles associated with it, or that it has one theoretical role comprising of the physicalist and dualist roles. The difference between the two is semantic. For what it’s worth, anecdotal evidence suggests that most favour the latter. I shall thereby proceed by assuming that the conditional analysis of ‘qualia’ entails that there is one theoretical role associated with the concept, but one that includes both the physicalist and dualist roles. Since there is room here to get talk of roles mixed-up, I shall proceed by using the terminology ‘functional roles’ to capture the individual reference-fixing roles, like the dualist and physicalist roles.¹⁶ These are to be contrasted with the (general) theoretical role of a concept, which can be seen to incorporate various distinct functional roles when our concept is given a conditional analysis.

Jackson (1998) offers something that looks very much like a conditional analysis of ‘freewill’, but not of ‘qualia’. Qualia, for him, like Lewis (1995), are analyzable with respect to the common-sense causal-cum-functional roles alone.¹⁷ The difference between Jackson and Braddon-Mitchell, in this regard, is a difference in how they partition the possible cases. Braddon-Mitchell supposes that there are some possibilities where the dualist roles help determine the reference of ‘qualia’, whereas the latter day Jackson, on account of his common-sense functionalism, seems to think otherwise. What this makes evident is that Jackson’s strategy for analyzing concepts doesn’t entail that ‘qualia’, or any other concept for that matter, ought to be given a conditional analysis. What it does, however, is prescribe conditional analyses where appropriate. That is, the partitioning of possible cases helps accommodate conditional theoretical roles; by suggesting when, in fact, there are distinct functional roles associated with the theoretical role of a given concept.

To clarify, any concept can be given a conditional analysis in the sense that we can analyze it with two or more conditionals. Nevertheless, for a lot of concepts, such an analysis would be superfluous. For instance, the relevant functional role for a ‘music player’ is that it plays music, which accounts for why gramophones, cassette players, tablets, mobile phones etc., all count as music players. Since there is one functional role, which is shared across all possible cases where we think the term applies, there is no need to give a conditional analysis that spells out distinct functional roles for the different possibilities, as part of its theoretical role. It will suffice to define its theoretical role in the following broad sense: *x* is whatever occupies the *x*-role. For other terms, however, such an analysis will be found wanting.

For example, if the proponents of the conditional analysis of ‘qualia’ are correct about their partitioning of the possibilities, our concept can’t be given a simple analysis, as we did with the notion of a music player. Our analysis of the concept, to the contrary, will consist of at least two conditionals, specifying two distinct functional

¹⁵ Alternative versions of the conditional analysis of phenomenal concepts include those provided by Hawthorne (2002) and Stalnaker (2002).

¹⁶ Note: ‘functional roles’ are to be used in Jackson’s neutral sense, where they are inclusive of non-causal descriptive roles.

¹⁷ This view is defended in Braddon-Mitchell and Jackson (1996/2007).

roles, akin to (CA1) and (CA2), as part of its theoretical role. The reason for this is that the functional roles that are relevant for fixing the reference of the term ‘qualia’ are context-dependent. More specifically, they depend on which worlds we take to be actual. As we saw earlier, if there actually are nonphysical states of the relevant type, which occupy the dualist roles, then it is these roles that fix the reference of ‘qualia’. If not, and the relevant physical roles are still actually occupied, then it is these physicalist roles that fix its reference.¹⁸ The same is not true of the term ‘music player’, as the functional roles relevant for fixing its reference don’t turn out to be context-dependent in this way. This difference between the two terms is discovered precisely by partitioning of the possibilities.

Where the method of partitioning tells us that conditional analyses are appropriate, the general form of such analyses will be something like the following:

(CA) If theory T1 is true, ‘t’ refers to whatever occupies the x1-role, if theory T2 is true, ‘t’ refers to whatever occupies the x2-role, ... and so on.

Analyses of this form are by no means problem-free.¹⁹ Nonetheless, they arguably pre-empt the problem of conceptual change. For instance, the typical dualist response to physicalist accounts of phenomenal consciousness, especially to ones like Australian materialism, is to claim that physicalists are simply eliminating the phenomenon in question, viz. qualia, and are using the same concept to refer to an altogether different set of entities. The conditional analysis of ‘qualia’ tells us the conditions under which these charges are justified. That is, to let the best physicalist theory fix its reference, *on the supposition that the actual world has the relevant spooky states*, is to eliminate the phenomenon in question and change our concept. However, to let these roles fix its reference, *on the supposition that the actual world lacks the relevant spooky states*, involves neither eliminativism nor conceptual change. This is because the context-sensitive nature of which roles are relevant to determining its reference, including the physicalist ones, are already specified in our initial analysis of the concept. The shift from a dualist theory to a physicalist one, then, needn’t involve conceptual change; at least, not to the extent that we are sensitive to the outcomes of partitioning of the possibilities, and analyse our concept accordingly.

Now recall the problem of conceptual change as it manifests itself with atoms. Dalton’s discredited theory arguably doesn’t share enough functional roles in common with its replacements, e.g. the theories of Thompson, Rutherford, Bohr etc. Were we to presuppose the RCL method—where there is a specific set of functional roles, which we take to be the theoretical role that determines the reference of a given concept—it follows that these different atomic theories have undergone conceptual change, and thereby are about different subject-matter.

Suppose, however, that the partitioning of possible cases where facts about atoms obtain, and ergo, cases where the term refers, tells us that if Dalton’s theory is true, atoms refer to whatever occupies the functional roles of his theory, whereas if Rutherford’s theory is true, it refers to whatever occupies the functional roles of this particular

¹⁸ For more detail on the context-sensitive nature of the roles, see [Majeed \(2015\)](#).

¹⁹ E.g. [Yetter-Chappell \(2013\)](#) argues that the conditional analysis of phenomenal concepts, which shares this structure, is circular.

theory, and so on. If this were the case, we can give a conditional analysis of the term ‘atom’, which shares the basic structure of (CA).

Moreover, we can argue that no conceptual change occurs because we factor in the possible cases into our analysis of the meaning of the theoretical term. That is, if atoms, by definition, are ‘whatever plays roles x_1 if Dalton’s theory is true, whatever plays roles x_2 if Rutherford’s theory is true, whatever plays roles x_3 if Bohr’s theory is true ...’²⁰, there is no sense in claiming that we have changed the concept if we abandon Dalton’s theory in favour of a replacement theory. Such a charge would only be justified provided we keep the term even if none of the theories of physics, which posit atoms and define their functional roles, turn out true. Which is to say we get conceptual change only if we keep the term even if the theoretical role—which subsumes all these functional roles—itself remains unrealized. But such (mis)use would, presumably, also go against what we learn from the partitioning of possible cases from which we learn what the theoretical role actually is. Both elimination and location, so it would appear, are sufficiently accounted for by the resultant conditional analysis.

5 Remaining problems

There is one major piece of housekeeping, which may be viewed as a significant problem depending on what you want from your conceptual analysis. For Jackson, conceptual analysis plays a modest role: it doesn’t help determine the fundamental nature of our world; rather, it helps determine what to say in less fundamental terms given an account of the world stated in more fundamental terms. Conversely, Canberra planners, as Nolan notes, typically employ conceptual analysis to play an immodest role: to help determine the fundamental nature of our world. In other words, the former concerns what to say about mysterious phenomena given certain suppositions about the actual world, including which features we take to be more or less natural-cum-fundamental. The latter project also (sometimes) involves attempts to demystify phenomena, but it additionally concerns which of these suppositions about the actual world hold true to begin with.

Conditional analyses can play the modest role, but not the immodest one. As we have seen, these analyses specify what the relevant functional roles are, and thereby help identify the referents as role-occupants—based on various considerations about how the actual world turns out. But since these roles themselves are dependent on which worlds we consider as actual, conditional analyses won’t help determine which worlds are, in fact, actual. For example, the conditional analysis of ‘qualia’ tells us if there actually are nonphysical states of the relevant type, ‘qualia’ refer to the occupants

²⁰ Note, “Dalton’s theory” etc. are used as shorthand for explicit statements of the relevant theories. The claim here, then, isn’t that “Dalton’s theory”, for instance, was written into the definition of the term ‘atom’ even before Dalton was born. The claim, rather, is that people were disposed all along to use the term ‘atom’ in the same way that it is used in Dalton’s theory — i.e. on the condition that the theory were presented to them as true. (This is what the partitioning of the possibilities is supposed to show). This is compatible with the theory never having been presented to them. It is, thereby, also compatible with them not knowing who Dalton is, or knowing how to explicitly define ‘atom’ in accordance with his theory.

of such-and-such dualist roles, whereas if such states are absent, they refer to the occupants of such-and-such physicalist roles. What this analysis doesn't do is tell us whether there are, in fact, nonphysical states of the relevant type. In this regard, it remains silent as to whether the dualist or physicalist is right about her ontology.

Construing this as a failure of the conditional analysis is to presume the immodest line. The conditional analysis of 'qualia', after all, wasn't meant to settle matters of ontology. Not directly anyway. Rather, it was intended to diffuse Chalmers's (1996) conceivability argument, and in a way that is compatible with our having the zombie intuition.²¹ Now, if modal arguments, like the conceivability argument, are our best reasons for dualism, and if we can undermine them by employing the conditional analysis, we can then, indirectly, use such an analysis to argue for physicalism. This is precisely what the original proponents of the analysis do. However, note that this involves additional moves, e.g. making independent claims about what best motivates dualism, and giving a high credence to the capability of a future neuroscience to account for all the phenomenal facts; moves that go way beyond the content of (CA1) and (CA2).

How does this bear on our ability to do serious metaphysics? This depends on how you interpret the location problem. On the one hand, the location problem is the problem of demonstrating precisely how to locate the mysterious facts in the non-mysterious ones—assuming the non-mysterious basic, fundamental facts are thus-and-so. Jackson's conceptual analysis, which prescribes conditional analyses where appropriate, can solve this problem. As we have seen, it tells us that were the fundamental facts thus-and-so, they entail facts about such-and-such mysterious phenomena.

On the other hand, the location problem is the problem of figuring out which mysterious facts are entailed by the non-mysterious fundamental ones *simpliciter*. This, I take, is what most metaphysicians understand by 'serious metaphysics' and take themselves to be doing. That is, they want to know what the actual world is like, including which facts are actually fundamental, as opposed to figuring out the various ways the world could turn out, based on various suppositions about what might constitute its fundamental facts. Jackson-style analyses won't help solve this problem given the end products of such analyses are always conditional: the non-mysterious facts *F* entail a particular set of mysterious facts *M*—but only on the proviso that the fundamental facts are thus-and-so.

These kinds of analyses are still useful, e.g. to undermine certain arguments. I think they are also illuminating, as they tell us which mysterious facts will be accounted for by a given fundamental theory of the world, and which ones will have to be abandoned. Nonetheless, they fall short when it comes to telling us which theory of the world is to be preferred. To the extent that Canberra planners are in the business of doing serious metaphysics understood in the second sense, they can't take their cues from Jackson, for to do so is to not recognize just how radically modest his account of conceptual analysis really is. And in light of this, they would also be wrong to look to this account to address the problems of conceptual change, which dog their attempts to do serious metaphysics.

²¹ See Chalmers (2005) and Alter (2007) for a critique, and Haukioja (2008) and Majeed (2014) for a defence.

In conclusion, the Canberra plan faces problems of conceptual change, which in turn inhibits its reliability in performing serious metaphysics. Jackson's conceptual analysis, by contrast, doesn't face these problems. However, it can't be of use for serious metaphysics either, as the machinery it has in place, which pre-empts the charges of conceptual change, renders it incapable of performing serious metaphysics. The take home slogan, then, is one of pessimism: 'no serious metaphysics from either a modest or immodest conceptual analysis!'²²

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Alter, T. (2007). On the conditional analysis of phenomenal concepts. *Philosophical Studies*, 134, 235–253.
- American Psychiatric Association. (1952). *Diagnostic and statistical manual of mental disorders* (1st ed.). Washington, DC: American Psychiatric Association.
- American Psychiatric Association. (1968). *Diagnostic and statistical manual of mental disorders* (2nd ed.). Washington, DC: American Psychiatric Association.
- American Psychiatric Association. (1980). *Diagnostic and statistical manual of mental disorders* (3rd ed.). Washington, DC: American Psychiatric Association.
- American Psychiatric Association. (2000). *Diagnostic and statistical manual of mental disorders* (4th ed.). Washington, DC: American Psychiatric Association.
- American Psychiatric Association. (2013). *Diagnostic and statistical manual of mental disorders* (5th ed.). Washington, DC: American Psychiatric Association.
- Baker, J. P. (2013). Autism at 70—Redrawing the boundaries. *New England Journal of Medicine*, 369, 1089–1091.
- Block, N., & Stalnaker, R. (1999). Conceptual analysis, dualism, and the explanatory gap. *Philosophical Review*, 108(1), 1–46.
- Braddon-Mitchell, D. (2003). Qualia and analytical conditionals. *Journal of Philosophy*, 100(3), 111–135.
- Braddon-Mitchell, D. (2005). The subsumption of reference. *British Journal for the Philosophy of Science*, 56, 157–178.
- Braddon-Mitchell, D., & Jackson, F. (1996/2007). *Philosophy of mind and cognition*. Blackwell.
- Carnap, R. (1963). Replies and systematic expositions. In P. Schilpp (Ed.), *The philosophy of Rudolf Carnap* (pp. 859–1013). La Salle, IL: Open Court.
- Chalmers, A. (2009). *The scientist's atom and the philosopher's stone: How science succeeded and philosophy failed to gain knowledge of atoms*. New York: Springer.
- Chalmers, D. J. (1996). *The conscious mind. In search of a fundamental theory*. Oxford: Oxford University Press.
- Chalmers, D. J. (2005) 'The two-dimensional argument against materialism'. Abridged version. In B. McLaughlin (Ed.), *Oxford handbook of the philosophy of mind*. Oxford: Oxford University Press (2009).
- Chalmers, D. J., & Jackson, F. (2001). Conceptual analysis and reductive explanation. *Philosophical Review*, 110(3), 315–361.
- Donaldson, T. (2015). Platitudes in mathematics. *Synthese*, 192, 1799–1820.
- First, M. B. (2012). The development of DSM-III from a historical/conceptual perspective. In K. S. Kendler & J. Parnas (Eds.), *Philosophical Issues in Psychiatry II: Nosology*. Oxford: Oxford University Press.
- Hacking, I. (1995). *Rewriting the soul: Multiple personality disorder and the sciences of memory*. Princeton: Princeton University Press.

²² I am grateful to Ben Blumson for comments on the draft. This paper was written as part of the project, 'New Directions in the Study of Mind', funded by the John Templeton Foundation.

- Haukioja, J. (2008). A defence of the conditional analysis of phenomenal concepts. *Philosophical Studies*, 139, 145–151.
- Hawthorne, J. (2002). Advice for physicalists. *Philosophical Studies*, 108, 17–52.
- Hey, P. R. (1934). Atoms. *The Scientific Monthly*, 38(6), 493–500.
- Jackson, F. (1994). Metaphysics by possible cases. *Monist*, 77, 93–110.
- Jackson, F. (1998). *From metaphysics to ethics*. Oxford: Oxford University Press.
- Johnston, M. (1992). How to speak of the colours. *Philosophical Studies*, 68, 221–263.
- Lewis, D. (1970). How to define theoretical terms. *Journal of Philosophy*, 67(13), 427–446.
- Lewis, D. (1972). Psychophysical and theoretical identifications. *Australasian Journal of Philosophy* 50: 249–258. Reprinted in Lewis, *Papers in Metaphysics and Epistemology*. Cambridge: Cambridge University Press (1999): 248–261.
- Lewis, D. (1994). Reduction of mind. In S. Guttenplan (Ed.), *Companion to the philosophy of mind*. Oxford: Blackwell.
- Lewis, D. (1995). Should a materialist believe in qualia? *Australasian Journal of Philosophy*, 73, 140–144.
- Majeed, R. (2014). A priori conditionals and the conceivability of zombies. *Philosophical Papers*, 43(2), 227–253.
- Majeed, R. (2015). Conceptual instability and the new epistemic possibility. *Erkenntnis*, 81(3), 613–627.
- Mayes, R., & Horwitz, A. V. (2005). DSM-III and the revolution in the classification of mental illness. *Journal of the History of Behavioral Sciences*, 41(3), 249–267.
- Nolan, D. (2005). *David Lewis*. Durham: Acumen Publishing.
- Nolan, D. (2009). Platitudes and metaphysics. In D. Braddon-Mitchell & R. Nola (Eds.), *Conceptual analysis and philosophical naturalism*. Cambridge: MIT Press.
- Papineau, D. (1996). Theory-dependent terms. *Philosophy of Science*, 63(1), 1–20.
- Ramsey, F (1931). *The foundations of mathematics and other essays*. In R. B. Braithwaite (Ed.), London: Routledge and Kegan Paul.
- Stalnaker, R. (2002). What is it like to be a zombie? In T. S. Gendler & J. Hawthorne (Eds.), *Conceivability and possibility*. Oxford: Oxford University Press.
- Stich, S. P. (1996). *Deconstructing the mind*. New York/Oxford: Oxford University Press.
- van Melsen, A. G. (1952). *From atomos to atoms: The history of the concept atom*. In H. J. Koren (trans.). Duplicates University Press.
- Yetter-Chappell, H. (2013). Circularity in the conditional analysis of phenomenal concepts. *Philosophical Studies*, 165(2), 553–572.