

Meanings and Methodologies

Justin C. Fisher – Southern Methodist University

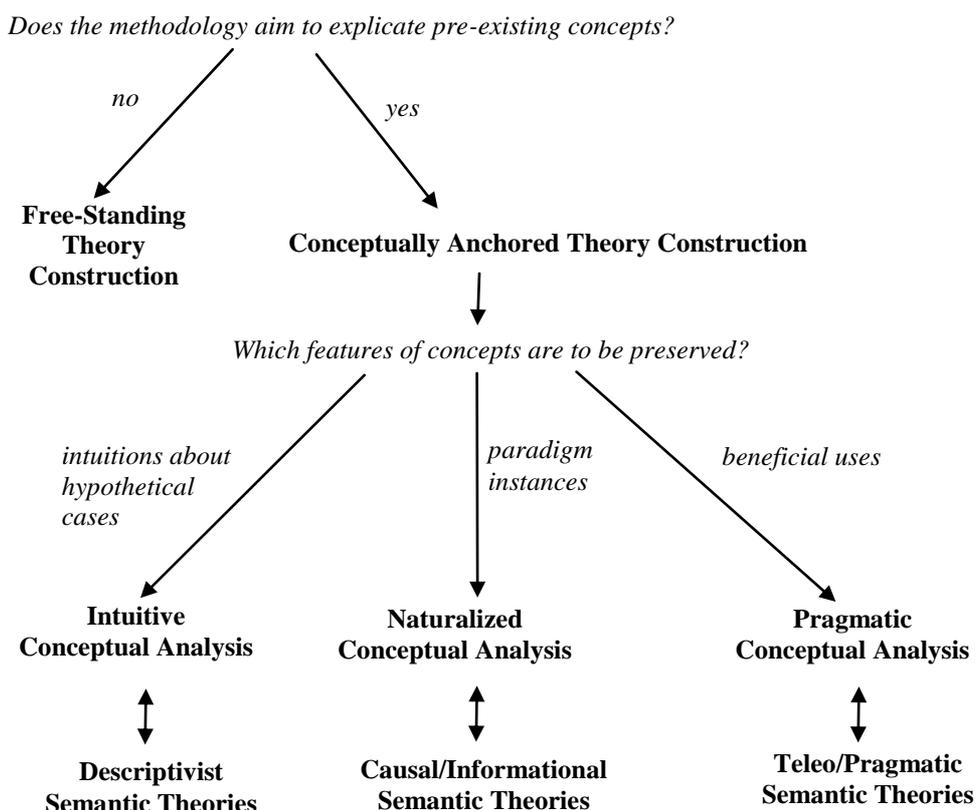
One new wave in the philosophy of mind involves connecting recent work in philosophy of mind to our ‘metaphilosophical’ understanding of methodologies for doing philosophy. This paper charts relations between (a) views in philosophy of mind and language regarding the correct application conditions, or ‘meanings’, of our words and concepts and (b) methodologies that people have proposed for doing philosophy, especially methodologies that have aimed to uncover the meanings of philosophical concepts like knowledge, freedom and justice. I identify three broad classes of theories of concept meaning. Two of these, descriptivist and causal/informational classes of theories, correspond closely to familiar philosophical methodologies – intuitive conceptual analysis and ‘naturalized’ analysis. A third, the teleo/pragmatic class, has many adherents in philosophy of mind but does not yet have a well-known corresponding philosophical methodology. To fill this gap, I describe a general methodology that I call Pragmatic Conceptual Analysis. I offer some examples of this methodology and argue that this methodology enjoys distinct advantages over more familiar philosophical methodologies.

I first lay out a space of possible methodologies and then link some of these to theories of concept meaning. Each of these methodologies is a guide to conceptual engineering, a guide that helps determine what concepts¹ to employ to determine exactly when and how we will apply these concepts and what beliefs to form employing these concepts. We might naturally think of each such methodology as involving two steps. In the first step, we articulate various desiderata regarding the conceptual framework we are attempting to engineer. Most plausible methodologies will agree on some very general desiderata – for example, that (*ceteris paribus*) our theories should be simple and self-consistent. But different methodologies may disagree upon how we might go about articulating further desiderata – for example, whether we should demand that proposed explications match pre-theoretic intuitions surrounding a concept. In the second step we then try to explicate the relevant concepts in a way that does well to meet the desiderata articulated in the first step. Different methodologies might propose different ways of resolving the conflicts that arise when different desiderata make incompatible demands.

There is a wide variety of possible conceptual engineering methodologies corresponding to the many possible desiderata we might impose upon our theory construction and to the many ways we might resolve tensions between conflicting desiderata. It is useful to categorize these competing methodologies on the basis of their answers to two questions (illustrated in Figure 4.1): *Does the methodology aim to explicate pre-existing concepts?* And if so, *which features of these concepts do we wish our explications to preserve?* We consider each question in turn.

¹ I think of concepts as mental particulars that play a folder-like information-coordinating role in cognition. In this, I follow a rich tradition in cognitive psychology and the philosophy of cognitive science. For a good introduction to this tradition, see Margolis and Lawrence (1999); for a view of concepts very similar to my own, see Millikan (1998, 2000). This tradition may be contrasted with an equally rich philosophical tradition that takes concepts to be abstract entities.

Figure 4.1. Methodologies for conceptual engineering.



1 Free-standing theory construction

We may begin with the question of whether a methodology's goal is to explicate pre-existing concepts. When a theorist engages in *conceptually anchored theory construction*, she aims to propose clear statements of application conditions which in some appropriate sense capture in a more rigorous way what certain pre-existing concepts were supposed to have captured all along. Conceptually anchored theory builders are constrained by the desideratum that their explications should preserve important features of pre-existing concepts. (In a moment we ask which features are worth preserving in this way.)

In contrast, the theorist who engages in *free-standing theory construction* may construct her theory from scratch and claim that as a whole, it is well worth considering, even if its various theoretical notions don't explicate any pre-existing concepts.^{2,3} Free-standing theory builders are constrained only by general desiderata

² My distinction between free-standing and conceptually anchored theory construction closely mirrors Clark Glymour's (2004) distinction between Euclidean and Socratic approaches.

³ It may be difficult to draw a principled line between what is conceptually anchored and what is free-standing. Even the most 'free-standing' of theory builders usually intends her theories to be linked in some way to pre-existing concepts, if only to those of observational evidence and simple logical relations. Such linkages may be taken to shed at least some light upon those pre-theoretical concepts. On the other side, most 'conceptually anchored' theorists are willing to wriggle free of pre-theoretic expectations when doing so yields enough theoretical usefulness. Even if there are some hard-to-call borderline cases, many are easy to call, and so this distinction is still a useful one.

regarding the conceptual frameworks they propose. For example, one might hope that free-standing theories will be simple, self-consistent, generally applicable, explanatorily powerful, predictively accurate and pragmatically useful. Different versions of free-standing theory construction might give these desiderata different weights, and some versions might impose additional desiderata.

Most philosophical work in the English-speaking ‘analytic’ tradition has been conceptually anchored. It has aimed to clarify what is supposed to have been captured by various ‘folk’ concepts – knowledge, justice, moral goodness, intentions, responsibility and the like. However, this tradition does contain some examples of free-standing theory construction. One clear example is Ruth Millikan’s (1984) introduction of theoretical concepts like ‘proper function’ to help account for how people and animals get around in their environments:

‘Proper function’ is intended as a technical term. It is of interest because it can be used to unravel certain problems, not because it does or doesn’t accord with common notions such as ‘purpose’ or the ordinary notion of ‘function’. My program is far removed from conceptual analysis; I need a term that will do a certain job, and so I must fashion one. (Millikan 1984, 18)

Other examples include eliminativism of folk psychology in favour of contemporary neuroscience (e.g., Churchland 1981), prescriptive Bayesian accounts in epistemology (e.g., Bovens and Hartmann 2003) and a great deal of work in pure logic.

There is a deep and useful connection between my distinction between conceptually anchored and free-standing theory construction and Thomas Kuhn’s (1962) distinction between ‘normal’ and ‘revolutionary’ scientific practice. For Kuhn, *normal science* is work that is anchored to (partly inchoate) concepts commonly employed within an accepted paradigm, while *scientific revolutions* are periods in which many theorists engage in free-standing theory construction, each attempting to propose a new paradigm to replace an existing one and to avoid the anomalies that plagued it.

Just as it is important that at least some people attempt revolutionary science, it is important that at least some philosophers engage in revolutionary or free-standing theory construction. For this approach sometimes produces good new ways of conceiving of problems – ways that may eventually take a seat within mainstream philosophical, scientific or even ‘folk’ conceptions of these problems. My commitment to the worthwhileness of free-standing theory construction is evidenced by the fact that I consider the present chapter (especially my articulation of Pragmatic Conceptual Analysis) to be largely an exercise in free-standing theory construction.

Unfortunately, free-standing theory construction is very difficult, precisely because it is so free of conceptual anchors to familiar and well-tested grounds. Free-standing theories usually must be very large and so usually offer only limited returns to people who aren’t prepared to immerse themselves completely in a new framework. In light of these drawbacks, there is much to be said for work that is conceptually anchored, to the concepts either of ordinary life or of some theory already quite familiar to at least some people. Conceptually anchored theory building is comparatively easy, yields comparatively high marginal returns and requires comparatively little departure from entrenched ways of conceiving of issues.

2 Conceptually anchored theory construction

By definition, conceptually anchored approaches aim to provide an important sort of continuity with pre-existing concepts. However, such approaches also typically demand that we depart *at least a little* from our antecedent understanding of our concepts, if only because such approaches seek explications that are *formal* and *explicit*, while pre-theoretic understandings are typically *informal* and *implicit*. But a proposed explication can't depart too far, lest it lose its claim to capture the same thing the pre-existing concept was supposed to capture. This opens a question: *which features of a pre-existing concept must a proposed explication retain if it is to count as a good explication of that concept?*

One tempting answer is to say that an explication should retain the *extension* of the explicated concept – the set of things to which that concept could be correctly applied. In what follows I often abbreviate ‘correct application conditions’ for a concept to the ‘meaning’ of that concept. However, it is worth emphasizing that this chapter focuses upon purely extensional aspects of concept meaning. Some people think there are finer-grained, or ‘intensional’, aspects of concept meaning – for example, aspects upon which HESPERUS and PHOSPHORUS differ in meaning. Such ‘intensional’ aspects are beyond the scope of this paper.⁴ It will be hard enough to determine how various methodologies relate to extensional aspects of meaning without worrying about further intensional aspects.

Unfortunately, there are at least two compelling reasons not just to say that the goal of explication is to preserve the extension or meaning of explicated concepts. First, it is controversial which theory of concept meaning is correct. Different semantic theories say that the meanings of our concepts depend on different factors, and hence these theories would identify different features of our concepts as the ones that would need to be preserved if our explications are to preserve concept meaning. Second, even after we adopt a particular theory of concept meaning, the question why we should *want* to seek explications that preserve this sort of meaning will remain. Might it not be the case that we should sometimes adopt *semantically revisionary explications* – explications that are worth adopting despite the fact that they require us to change what some of our concepts mean?

We can forestall these worries by considering specific proposals regarding which features should be preserved in explicating a concept. Then we can ask which (if any) of these specific proposals preserve the features that are determinative of concept meaning. (Our answer will depend, of course, upon which semantic theory we think is correct.) And we can ask which (if any) of these specific proposals provide explications that are worth adopting, even if (at least according to some semantic theories) doing so requires semantic revision. I give special attention to three specific proposals:

⁴ For my own views on ‘intensional’ aspects of meaning, see Fisher (2006, n.d. 2, esp. §§1.2, 7.6).

1. *Pragmatic Conceptual Analysis* aims to preserve the ways in which usage of the concept in question has regularly delivered benefits.
2. *Intuitive Conceptual Analysis* aims to preserve the truth of various intuitions surrounding that concept.
3. *Naturalized Conceptual Analysis* aims to preserve what we would count as paradigm instances of that concept.

(Readers sometimes balk at my using the label ‘conceptual analysis’ for anything other than Intuitive Conceptual Analysis. I use this label because these conceptually anchored methodologies all share the general goal of traditional conceptual analysis: carefully exploring the relevant features of pre-existing concepts to state the application conditions of those concepts in other terms. Admittedly, these methodologies don’t involve ‘analysis’ in the old-fashioned sense of breaking concepts apart into component parts. But virtually everyone is now convinced by psychological research⁵ that few, if any, concepts have the neat definitional structure that would be required for them to be ‘broken apart’ in this way. So even contemporary practitioners of Intuitive Conceptual Analysis think ‘analysis’ involves, not breaking concepts into component parts, but instead examining the roles concepts play in broader cognition. Naturalized and Pragmatic Conceptual Analysis hold that the scope of analysis – including the roles concepts play in interaction with external objects – is a bit more broad, but that doesn’t make this examination any less of an analysis.)

These three answers embody three general methodologies for conceptual engineering. But they also correspond to three general theories of concept meaning, where each methodology yields explications that preserve the corresponding sort of concept meaning. In what follows, I continue to take *methodologies* as the primary locus of discussion. As I discuss each methodology, I note the links between it and the corresponding theories of concept meaning (these links were also noted at the bottom of Figure 4.1).

In considering these three methodologies (and the corresponding theories of concept meaning), it will help to also consider how each might apply to a particular example: the familiar philosophical debate regarding whether *free action* is compatible with determinism.

3 Pragmatic Conceptual Analysis

I begin with the methodology I call Pragmatic Conceptual Analysis, as I think it enjoys important advantages over the more familiar methodologies considered below. Pragmatic Conceptual Analysis proposes that we ‘reverse-engineer’ an existing conceptual scheme to determine how it works as well as it does so that we might then modify it to more consistently deliver benefits in these ways. Hence, the key desideratum imposed by Pragmatic Conceptual Analysis is that our explications preserve the ways in which our applications of pre-existing concepts have regularly delivered benefits.

⁵ For an overview of relevant research, see Laurence and Margolis (1999).

While Pragmatic Conceptual Analysis has not yet received as much attention as the other methodologies I consider, there have been a few articulations of this general sort of methodology.⁶ One quite good articulation is given by the epistemologist Sally Haslanger:

[T]he best way of going about a project of normative epistemology is first to consider what the point is in having a concept of knowledge: what work does it, or (better) could it, do for us? And second, to consider what concept would best accomplish this work. (Haslanger 1999, 467)

I say a great deal elsewhere about the various particular ways in which one could and should develop this sort of methodology. But we can get a good sense of the general way in which Pragmatic Conceptual Analysis works without going into excessive detail.

As an example let us consider a classic philosophical question: *is free action compatible with determinism?* The answer to this question depends, of course, upon which possible actions are to be counted as falling under our concept FREE ACTION. This is something that Pragmatic Conceptual Analysis can help determine.

As with other methodologies, Pragmatic Conceptual Analysis may naturally be divided into two steps. In the first step the desiderata that will constrain our choice of explications are articulated. For Pragmatic Conceptual Analysis, the primary desideratum embodies a *job description* that outlines the regular ways in which our use of a concept has regularly delivered benefits. For present purposes, we may define ‘benefit’ as anything the person using the concept in question has practical reason to pursue.⁷ I leave it to other philosophers to determine exactly what the pursuitworthy benefits are, but I presume that there are at least some uncontroversial examples of these, including many instances of achieving happiness or satisfaction and many instances of avoiding pain, injury or death. It is an empirical question how our concepts have delivered such benefits, and hence it is an empirical question what job description will be delivered by a reverse engineering analysis of our use of a shared concept.

For the purpose of illustration, I stipulate an answer to this question in the case of FREE ACTION (an answer that is a fairly plausible empirical hypothesis anyway), and we may consider what explication to adopt if this stipulation turns out to be true.⁸ Let us suppose that it will turn out that *determinism* is true and hence that each human action has a complete cause in the distant past. But let us also suppose that it will turn out that there are systematic patterns in what is caused by our categorization of some actions as ‘free’ and others as ‘not free’. In particular, let us suppose that the regular way in which categorizing actions as ‘free’ has yielded benefits is by leading us to reward or punish the people who perform such actions, which in turn has encouraged people with well-functioning deliberative systems to perform beneficial actions and discouraged them from performing

⁶ Closely related methodologies have been proposed in epistemology by Edward Craig (1990) and in philosophy of science by James Woodward (2004). I offer a much more detailed articulation and defence of Pragmatic Conceptual Analysis in Fisher (2006, n.d. 2) and also survey the relations between it and other somewhat similar approaches, including William James (1906) and Rudolf Carnap (1950).

⁷ Elsewhere I consider various detailed understandings of ‘benefit’ that might be employed by different versions of Pragmatic Conceptual Analysis and the relative merits of each. See Fisher (2006, n.d. 2).

⁸ In other works, I consider ways in which research by experimental philosophers (and other experimentalists) might provide a better-justified job description for this concept; see Fisher (forthcoming, n.d. 2, ch. 5).

harmful actions. And let us suppose that the regular way in which categorizing actions as ‘not free’ has yielded benefits has been by getting us to seclude and give medical treatment to people who have a tendency to perform harmful actions not under the control of well-functioning deliberative systems.

This pattern of benefit delivery embodies a job description that one might reasonably hope any good explication of FREE ACTION would be able to perform. To allow us to continue reaping benefits in these ways, an explication must count as ‘free’ those sorts of actions whose performance will be well regulated by standard practices of offering praise and blame (or punishments and rewards), and it must categorize as ‘not free’ those actions whose performance is more effectively regulated by seclusion and medical treatment.

The second step of Pragmatic Conceptual Analysis takes a job description like this and seeks to find an explication that will do optimally well at fulfilling it. Since we are presuming that determinism is true, this notion will need to be *compatibilist* – it must count some actions as free even if they have complete causes in the distant past. This compatibilist conclusion does not depend much upon the particular ways in which use of FREE ACTION delivers benefits. So long as determinism is true (or close enough to true) and categorizations of actions as ‘free’ and ‘not free’ are beneficial in regular ways, only compatibilist explications of FREE ACTION will allow us to continue to achieve benefits in these ways.

Given particular empirical details about how such categorizations yield benefits, Pragmatic Conceptual Analysis can lead to adoption of a *particular sort* of compatibilism. For example, given my empirical stipulations above, the notion that will best fulfil the identified job description might be something like ‘action susceptible to causal control by the normal human deliberative processes that normally are sensitive to considerations of rewards and benefits.’

We may define the *pragmatic meaning* of a concept as the explication that Pragmatic Conceptual Analysis would deliver for that concept. (Since there are different ways of working out the details of Pragmatic Conceptual Analysis, there will likewise be different but closely related technical notions of pragmatic meaning.) Pragmatic meaning is closely related to *teleo/pragmatic* semantic theories that have been defended by various authors, including William James (1906), Frank Ramsey (1927), Ruth Millikan (1984), Anthony Appiah (1986), David Papineau (1987), Fred Dretske (1988), Jamie Whyte (1990) and Simon Blackburn (2005). In addition, I argue elsewhere (Fisher 2006; n.d. 2, chs. 4–6) that we should accept success-linked theories of content in general and that, in particular, we should think of pragmatic meaning as capturing ‘the meanings’ of our concepts.

Note too how worthwhile it is to use Pragmatic Conceptual Analysis to identify the pragmatic meaning of a concept. The key desideratum placed by Pragmatic Conceptual Analysis is that we find explications that will deliver pursuitworthy benefits if employed. Because of this, there is a guarantee that Pragmatic Conceptual Analysis will deliver explications worth employing.⁹ While some people may resist thinking that *the* meaning of

⁹ This guarantee is subject to several caveats. E.g., it may be that some concepts – including perhaps PHLOGISTON – just don’t merit continued use, even after we use Pragmatic Conceptual Analysis to tweak them. When a concept turns out not to be all that useful, it may be reasonable to abandon it and adopt another conceptual framework – e.g., that of modern chemistry – in its place. Such caveats do not weigh against the

a given concept is its pragmatic meaning, these people still must acknowledge that there is a strong case for stipulating that, henceforward, we will use a given concept in accord with its pragmatic meaning.¹⁰

My deepest objection to competing conceptually anchored methodologies is that they offer no similar guarantees of beneficial results. I drive home this general concern in Section 7 below. But first let us see what these other methodologies are.

4 Intuitive Conceptual Analysis

I contrast Pragmatic Conceptual Analysis with two other conceptually anchored methodologies. I call the first *Intuitive Conceptual Analysis*. This familiar philosophical methodology aims to deliver explications that preserve the truth of intuitions surrounding the concepts being explicated. One staunch defender of this methodology is Frank Jackson:

[H]ow should we identify our ordinary conception? The only possible answer, I think, is by appeal to what seems to us to be most obvious and central about free action, determinism, belief, or whatever, as revealed by our intuitions about possible cases. [...] Intuitions about how various cases, including various merely possible cases, are correctly described in terms of free action [...] are precisely what reveal our ordinary conceptions of free action. (Jackson 1998, 31)

There are different versions of Intuitive Conceptual Analysis corresponding to the different ways in which one might use intuitions to constrain one's choice of explication.

One potential point of difference involves the sort of intuitions to be appealed to. In the quote above, Jackson calls upon intuitions about whether a concept would apply in various possible cases. But other versions of Intuitive Conceptual Analysis might also give weight to intuitions as to which general claims should hold true regarding a concept. (For example, if we have the strong intuition that all bachelors are men, one might take this generalization itself to be something worth preserving in our explication of BACHELOR.) And some versions of Intuitive Conceptual Analysis might give special weight to strong intuitions that certain things definitely are instances of a given concept. (For example, one might find it intuitively obvious that this rock is solid and might take preserving the truth of this intuition as an important desideratum for any analysis of SOLIDITY.)

A second potential point of difference involves the question of how much weight to give to the intuitions of different people. Some versions of Intuitive Conceptual Analysis would favour the 'pre-theoretic' intuitions of ordinary people not in the grip of any theory, while other versions might favour the intuitions of people who are empirically well informed, who have significant practical experience or who have spent years reflecting carefully upon a topic.

claim that Pragmatic Conceptual Analysis generally yields explications worth adopting and generally does markedly better at this than rival methodologies.

¹⁰ I expand on this idea in Fisher (2006, ch. 4, n.d. 1, n.d. 2, ch. 4), where I argue that Pragmatic Conceptual Analysis has 'normative authority' in that we generally have both *practical* and *epistemic* reason to embrace the explications it delivers.

Third, different versions of Intuitive Conceptual Analysis might endorse different ways of resolving conflicts between intuitions. For example, some versions might allow us to disregard intuitions that can be ‘explained away’ as stemming from unreliable sources. Some might also propose ways of giving different amounts of weight to intuitions that are felt to have different degrees of ‘strength’ or confidence.

Intuitive Conceptual Analysis is a mainstay of recent analytic philosophy.¹¹ This is evidenced by the fantastic menagerie of ‘intuition pumps’ with which all analytic philosophers are painfully familiar, including trolley problems, Gettier cases, the Chinese Room, Swampman and Twin Earth.

Intuitive Conceptual Analysis is also closely related to traditional *descriptivist theories of meaning*, which hold that the meaning of a person’s concept is determined by the description(s) that that person (perhaps tacitly) associates with that concept (Frege 1892, Russell 1905, Strawson 1950). Insofar as we expect intuitions to do a good job of revealing tacit descriptions, we can expect that some version of Intuitive Conceptual Analysis would do well to reveal meanings, in that term’s descriptivist sense.

Most people working in philosophy of mind and language have abandoned descriptivism and moved on to either causal/informational or teleo/pragmatic theories, according to which meaning is determined in large part by patterns of causal interaction between us and things in the world.¹² The mass exodus from descriptivism bodes poorly for the hope that Intuitive Conceptual Analysis will be a good way to reveal the meanings of concepts. Granted, some intuitions may track the semantically relevant external causal relations. Indeed, insofar as our intuitions are shaped by practical experience (via either learning or evolution), it may be especially likely that we will end up having quite good intuitions regarding ordinary practically relevant cases. However, Intuitive Conceptual Analysis has been characterized by its consideration of cases that are far from ordinary: Swampman, the Chinese Room, TrueTemp, Frankfurt’s neuroscientist. For such extraordinary cases, there is no reason to think intuitions will do especially well at tracking truths about the referents we have latched onto via the sorts of causal interaction that are deemed semantically relevant by most contemporary semantic theories.

Furthermore we now know that folk intuitions *often do fail* when extrapolated to extraordinary cases. Despite intuitions to the contrary, solid objects aren’t impermeable at small scales, water isn’t an indivisible element, species don’t have immutable essences, and space-time isn’t Cartesian at small scales or high speeds. It would take great hubris to expect folk philosophical intuitions to fare better than folk intuitions surrounding other concepts. Descriptivism offered one potential basis for such hubris: intuitions would be trustworthy if they reflected tacit descriptions that determined reference. But once we reject descriptivism, we lose grounds for optimism about the reliability of intuitions regarding extraordinary cases, and hence we lose grounds for trusting Intuitive Conceptual Analysis to correctly reveal the meanings of our concepts.

¹¹ For an alternative point of view, see Capellen (2012), which argues that philosophers have relied upon intuitions much less often than many have thought. I remain unconvinced but will not argue the point here. The primary purpose here is just to consider potential advantages and disadvantages of an intuition-driven methodology, regardless of how many people have actually used it.

¹² The exodus from descriptivism was spurred in part by classic examples by Kripke (1972), Putnam (1973) and Burge (1979). These include examples where people don’t have in mind descriptions sufficient to pick out their referents uniquely (e.g., Kripke’s Feynman case, Putnam’s Twin Earth), and examples where the descriptions people do have in mind don’t even fit their actual referents (e.g., Kripke’s Gödel/Schmidt case, Burge’s arthritis case).

Even while our best semantic theories give us strong reason to doubt our intuitions, a great deal of work in analytic philosophy still proceeds by Intuitive Conceptual Analysis. This is an unstable situation. Sooner or later some sort of correction will need to be made to resolve it: (a) philosophers of mind and language will need to reincarnate something much like descriptivism,¹³ or (b) analytic philosophers can keep pumping intuitions but give up their goal of revealing the correct application conditions of philosophical concepts, or else (c) we will need to shift to methodologies that better fit our state-of-the-art semantic theories.

Further challenges for Intuitive Conceptual Analysis will become apparent when we consider how this methodology would apply to the example of FREE ACTION. Frank Jackson, a leading defender of Intuitive Conceptual Analysis, consults his own intuitions surrounding FREE ACTION and concludes,

Speaking for my part, my pre-analytic conception of free action is one that clashes with determinism. (Jackson 1998, 44)

One might wonder whether Jackson's incompatibilist intuitions actually match the relevant intuitions of whoever the people are whose intuitions would be relevant to certain versions of Intuitive Conceptual Analysis.¹⁴ However, let us suppose for the sake of illustration that Jackson has applied Intuitive Conceptual Analysis correctly. Even if many real people's intuitions differ from Jackson's, they might just as easily have been the same as his, and this possibility is enough to illustrate how Intuitive Conceptual Analysis may fail to be useful.

Given that determinism is true (or close enough to true), Jackson's incompatibilist analysis of FREE ACTION does indeed fail to be useful, for it lumps all our actions together in a single category: *unfree*. This fails to offer many of the useful distinctions that we ordinarily use the concept FREE ACTION to make – for example, the distinction between the acts of ordinary vandals and the acts of sleepwalkers.

Jackson, himself recognizing this failing of his incompatibilist analysis, writes:

[It is] only sensible to seek a different but 'nearby' conception that does, or does near enough, the job we give [to the concept being analyzed] in governing what we care about, our personal relations, our social institutions of reward and punishment, and the like, and which is realized in the world. (Jackson 1998, 45)

Here Jackson admits that Intuitive Conceptual Analysis may lead us to explications that aren't very useful and that when this happens, it is 'only sensible' to employ a fallback strategy instead. Jackson's fallback strategy is to seek an explication capable of doing the jobs we give that concept in governing the things we care about. This is a fairly good characterization of Pragmatic Conceptual Analysis. Hence, Jackson admits that when push

¹³ Some contemporary philosophers aim to do just that, including Lewis (1984) and Chalmers (2002, 2004).

¹⁴ For example, many experimental subjects have intuitions that disagree with Jackson's (Nahmias et al. 2005, 2006). I discuss the significance of this research in Fisher (forthcoming, n.d. 2, ch. 5).

comes to shove – when Intuitive Conceptual Analysis and Pragmatic Conceptual Analysis disagree – it is ‘only sensible’ to embrace the guidance of Pragmatic Conceptual Analysis.^{15, 16}

This point is really quite general. We saw above that (given a few reasonable caveats) Pragmatic Conceptual Analysis guarantees that its explications will be useful. This is a guarantee that Intuitive Conceptual Analysis cannot match. Our intuitions have been shaped by (individual, cultural, and evolutionary) experience in various situations. These experiences have probably tended to shape our intuitions so that they fit *fairly well* with what has been beneficial in common cases and/or important ones. But our intuitions still must be a fallible guide to what is beneficial even in these cases; they become very fallible indeed as we attempt to apply them to cases that differ significantly from the ordinary high-stakes cases they have been most strongly shaped to deal with.

Since our intuitions probably correlate somewhat well with the ways in which our concepts are actually useful, our intuitions are a reasonable starting point for generating plausible hypotheses regarding what sorts of useful work our concepts are doing. But additional empirical evidence often reveals that our intuitions have overlooked some ways in which our concepts have regularly delivered benefits and that they are significantly mistaken about others. When faced with such discoveries, Pragmatic Conceptual Analysis and Intuitive Conceptual Analysis point us in opposite directions. Intuitive Conceptual Analysis proposes that we enshrine intuitions that we have discovered to be poor guides to what is beneficial, while Pragmatic Conceptual Analysis proposes that we abandon these misguided intuitions and instead continue to reap benefits in the regular ways we have been reaping them.

At such a point, Intuitive Conceptual Analysis calls upon us to forgo the sorts of benefits that our use of a concept has regularly delivered. What does it offer in exchange? Only the preservation of pre-existing intuitions that we now know fail to track what is beneficial. When faced with a choice between achieving tangible benefits and catering to intuitive prejudice, there is strong reason to follow the tangible benefits. As Jackson himself admitted, when Pragmatic Conceptual Analysis and Intuitive Conceptual Analysis diverge, it is ‘only sensible’ to follow Pragmatic Conceptual Analysis.

I have just argued that we should not use our intuitions surrounding a concept as the final arbiter in determining what explication of that concept to adopt. But it is worth noting that there are at least three limited ways in which intuitions might play a central role in philosophical analysis, even on my view.

¹⁵ To be fair, Jackson insists that this application of Pragmatic Conceptual Analysis will be a *revisionary* departure from the concept’s original meaning. Whether Jackson is right depends upon which semantic theory is correct. I have argued elsewhere (Fisher 2006, ch. 4, n.d. 1, n.d. 2, ch. 4) that we should embrace a teleo-semantic theory upon which Pragmatic Conceptual Analysis is semantically conservative, not revisionary.

¹⁶ Even while ‘sensibly’ embracing a compatibilist explication of FREE WILL for ordinary usage, one might still want to coin other concepts (e.g., perhaps LIBERTARIAN FREE WILL) for other more esoteric uses. I encourage introducing as many technical concepts as you like into the marketplace of ideas and seeing how they fare. Concepts explicated by Pragmatic Conceptual Analysis will generally be robust competitors in this marketplace as they will have been honed to continue yielding pursuitworthy benefits. But it’s possible that some alternative concepts without this promising pedigree might find a way to earn their keep as well. Indeed, I’ve already acknowledged this as a worthy goal of free-standing theory construction.

First, as just noted, our intuitions surrounding a concept may serve as a good source for initial hypotheses regarding the usefulness of that concept, hence as a good starting point for Pragmatic Conceptual Analysis. But they are just a starting point. When we discover that our intuitions are mistaken regarding the useful work a concept has been doing, then, on my view, we no longer have any reason to allow these intuitions to continue to constrain our theorization.

Second, if we are to have reason to embrace the explications produced by Pragmatic Conceptual Analysis, we will want to use a version of this methodology which is defined in terms of a notion of benefit that we have reason to pursue. For all I know, it may be that intuitions about what sorts of things are worth pursuing will play a large role in helping to choose this notion of pursuitworthy benefit. But notice that these are intuitions about what consequences are worth pursuing rather than about the concept(s) being analysed (e.g., FREE ACTION). This fact sets this approach apart from Intuitive Conceptual Analysis.

Third, Pragmatic Conceptual Analysis draws upon empirical accounts of how various concepts have regularly delivered benefits. In producing such empirical accounts, we might draw upon the sorts of intuitions scientists use to determine which empirical generalizations are supported by empirical evidence. These are quite different from the intuitions called upon by familiar forms of Intuitive Conceptual Analysis.

Hence, I stake out a fairly moderate position regarding the use of intuitions in philosophy and allow for at least three limited ways in which intuitions might be quite useful. However, I maintain that Intuitive Conceptual Analysis is at best a starting point and that strong empirical evidence about the actual usefulness of our concepts should trump our intuitions surrounding those concepts. It is quite reasonable to suspect that contemporary analytic philosophers have pushed our intuitions far beyond the limits of their usefulness. Rather than continue to pump intuitions about fantasy swampmen, esoteric trolley problems or strange new variants of the Gettier problem, philosophers would do well instead to seek good empirical evidence regarding the sorts of useful work that our concepts have actually been doing for us and then to seek explications that would allow our concepts to do this work more efficiently.

5 Naturalized Conceptual Analysis

Let us move on to consider another conceptually anchored methodology, one that takes as its conceptual anchor the various ‘paradigm instances’ of a given concept. The job of the philosopher or scientist is to go out and discover empirical facts about what (if anything) unifies these paradigm instances as a single sort of natural phenomenon. One strong advocate of this methodology is Hilary Kornblith:

We begin, often enough, with obvious cases, even if we do not yet understand what provides the theoretical unity to the kind we wish to examine. *Understanding what that theoretical unity is* is the object of our study, and it is to be found by careful examination of the phenomenon, that is, something outside of us, not our concept of the phenomenon, something inside of us. (Kornblith 2002, 10–11)

I call this methodology *Naturalized Conceptual Analysis*. Philosophers commonly call such approaches ‘naturalized’ because they draw heavily upon empirical findings from natural sciences. Such approaches are

also ‘conceptual analyses’ in that they take pre-existing concepts as their launching point and aim to arrive at good explications of those concepts.

Advocates of Naturalized Conceptual Analysis take this methodology to be a mainstay of empirical science,¹⁷ which clearly has delivered useful analyses for ordinary concepts like FISH, WATER, PLANET, and LIGHTNING. Naturalized Conceptual Analysis also enjoys a growing following in philosophy, including proposed analyses of KNOWLEDGE (Quine 1969; Kornblith 2002), EMOTION (Griffiths 1997), COLOUR (e.g., Hilbert 1987), CONSCIOUSNESS (Dennett 1991) and MORAL GOODNESS (e.g., Boyd 1988).

Naturalized Conceptual Analysis pairs naturally with *causal* and *informational* theories of reference, like those proposed by Saul Kripke (1972), Gareth Evans (1973), Richard Boyd (1988), Jerry Fodor (1990) and Robert Rupert (1999). Corresponding to the many different particular ways that one might spell out such theories of reference, there are different versions of Naturalized Conceptual Analysis.

Different versions of Naturalized Conceptual Analysis might disagree about what makes something count as a paradigm instance of a concept. In the passage above, Kornblith suggests that we start with ‘obvious’ cases. This suggests that the paradigm instances are those which intuitively strike us as especially clear instances of a given concept.¹⁸ Other proponents of Naturalized Conceptual Analysis take the paradigm instances to be things that were present at the initial baptism of a concept or term (Kripke 1972). Others take paradigm instances to be the things that are a source of information we associate with the concept (Evans 1973, Boyd 1988). Still others take paradigm instances to be the things which cause us to token the concept most consistently (Rupert 1999) or most robustly (Fodor 1990).

Different versions of Naturalized Conceptual Analysis might also disagree about what sort of theoretical unity to seek among the paradigm instances. For any given set of paradigm instances, there will likely be numerous more-or-less natural kinds which unify most of the instances. For example, let’s grant for the sake of argument that our shared concept FREE ACTION somehow designates a set of paradigm instances. The problem is that these instances will have *all kinds* of features in common. How will we know which sort of commonality to look for? We need some clear specification of what sort of commonality is relevant even to get in the ballpark of a good analysis. (Devitt and Sterelny [1987] call this the ‘qua problem’.)

There are several potential ways to deal with this problem. Adding a set of paradigm counterinstances might rule out many candidate ways of unifying the paradigm instances. Counting some ways of unifying paradigm instances as ‘more natural’ than others (cf. Lewis 1983) might give us a ‘natural’ way to prefer some candidates to others. Or we might prefer those (‘natural’) commonalities that include minimally many further instances beyond the given paradigm instances. However, none of these proposals seems to capture the plausible fact that on the basis of a given set of paradigm instances, a person might form a concept of *water* (i.e., the chemical compound H₂O) or a concept of *ice* (i.e., the solid phase of H₂O) or might even form both these

¹⁷ I think that we would actually do better to consider much normal scientific practice as involving Pragmatic Conceptual Analysis rather than Naturalized Conceptual Analysis, but I leave those arguments for elsewhere.

¹⁸ Such an appeal to intuition would blur the line between Naturalized Conceptual Analysis and Intuitive Conceptual Analysis, as many proponents of sophisticated versions of Intuitive Conceptual Analysis (e.g., Jackson 1999; Chalmers 2002) are inclined to do.

concepts. Perhaps for this reason, many advocates of Naturalized Conceptual Analysis have taken it that a given concept provides conceptual anchors not only to a set of paradigm instances but also to a sortal that specifies what sort or kind these instances are all supposed to belong to. Two concepts (like WATER and ICE) with the same paradigm instances might refer to different natural kinds in virtue of the fact that the two concepts are somehow associated with different sortals.

The need to associate concepts with sortals presses the advocate of Naturalized Conceptual Analysis to lean upon one of the above methodologies. Should we determine the sortal by consulting someone's intuitions surrounding the given concept? This move would lead to many of the disadvantages of Intuitive Conceptual Analysis mentioned above. Or should the sortals be determined in some other way – perhaps even by something like Pragmatic Conceptual Analysis? Doing this would help to ensure that Naturalized Conceptual Analysis will be useful but at the cost of conceding that Pragmatic Conceptual Analysis is a good guide after all.

There is a deeper problem here, paralleling the deep problem recognized above in Intuitive Conceptual Analysis. Even when we can find a way to carry Naturalized Conceptual Analysis through to a determinate conclusion, there is no guarantee that this conclusion will be all that useful – there is no guarantee that the analysis will latch on to a natural kind that it is worthwhile to use a shared concept to track. For example, a great deal of work in epistemology has taken the paradigm instances of KNOWLEDGE to be abstract cases like 'I think I therefore I am' and ' $2 + 3 = 5$ '. But there is no guarantee that the best account of what these paradigm cases have in common will be at all useful when applied to our ordinary beliefs regarding important empirical matters.

Pragmatic Conceptual Analysis does give us a guarantee like this: if a shared concept delivers regular benefits that make it worth having, then Pragmatic Conceptual Analysis will *at least* preserve this usefulness and often extend it. By contrast, Naturalized Conceptual Analysis might sometimes take what had been a quite useful concept and have us focus the application of that concept upon some natural kind that it is not nearly so useful for us to track.

Jackson's comments about what is 'sensible' apply here, too. When push comes to shove – when one is forced to choose between preserving beneficial usage and preserving paradigm instances – it is sensible to abandon the so-called paradigm instances and instead to accept the guidance of Pragmatic Conceptual Analysis.

6 Other methodologies

I have considered four general conceptual engineering strategies. Free-standing theory construction aims only to produce a useful conceptual framework, and is not concerned about explicating pre-existing concepts. Pragmatic Conceptual Analysis aims to preserve beneficial uses of our pre-existing concepts. Intuitive Conceptual Analysis aims to preserve intuitions surrounding our concepts. Naturalized Conceptual Analysis aims to preserve the paradigm cases and sortals associated with our concepts.

Some philosophical projects clearly employ just one of these methodologies. But more often, philosophers employ some combination of them. Conceptually anchored projects often introduce some free-standing technical terminology to help articulate their analyses. Naturalized approaches often lean heavily upon

intuitive approaches to specify sortals. Sophisticated advocates of intuition-driven approaches often allow that some intuitions might licence empirical findings to supplant other intuitions (Lewis 1984; Jackson 1998; Chalmers 2002); some advocates are also willing to turn to pragmatic approaches when intuitions fail to be useful (e.g., Jackson 1998; Schmidtz 2006).

In addition to the desiderata placed by these four methodologies and their hybrids, there might be other desiderata that one might think conceptual engineers should heed. In particular, there might be other aspects of pre-existing concepts that one might hope a proposed conceptual analysis would preserve. I briefly present the two most plausible candidates I have encountered.

First, some recent philosophical arguments have attempted to use various subtle linguistic findings to help guide selection of philosophical analyses. For example, Jason Stanley (2005) argues that the word ‘know’ fails various subtle linguistic tests for context sensitivity. [[*]]Stanley takes this linguistic evidence to militate against accepting a contextualist analysis of KNOWLEDGE.

Second, people are often quite reluctant to accept a proposed explication which fails to preserve the metaphors they had previously associated with a concept (e.g., Niyogi 2005). For example, we now know that *oxidizable materials* are actually responsible for much of what *phlogiston* was once thought to do (e.g., enabling burning and rusting). However, phlogiston theorists were strongly attached to the metaphor of phlogiston as a fundamental sort of fluid released from objects when they burn or rust. In contrast, oxidizable materials encompass a diverse assortment of substances that are not *released* through oxidation but are instead *transformed* into oxidized materials. Our contemporary understanding of oxidizable materials fits quite poorly with the old metaphors for phlogiston, and this might explain some people’s reluctance to embrace oxidizable material as an explication of PHLOGISTON.

One might take it as a stand-alone desideratum that conceptual analyses should deliver explications that preserve subtle linguistic features of analysed concepts or that allow us to continue using familiar metaphors for our concepts. However, my own inclination is to grant that like intuitions and paradigm cases, linguistic cues and associated metaphors are a mere source of defeasible evidence regarding the ways in which a concept is successfully used. I have no problem with using any of these as sources for initial empirical hypotheses regarding the usefulness of concepts. However, once we get strong empirical evidence that a concept has regularly been delivering benefits in some other way than what these cues initially suggest, we will have strong reason to prefer explications that preserve these beneficial uses to explications that sacrifice benefits to preserve these other cues. When we discover that our metaphors or our linguistic practices are in tension with the continued beneficial usage of our concepts, we should abandon the old metaphors or linguistic practices, not the beneficial uses.

Similar considerations apply regarding *any* conceptual anchor one might dream up. Suppose we are considering a version of Pragmatic Conceptual Analysis defined in terms of a notion of benefits that we have reason to pursue and suppose we are comparing this against some other conceptually anchored methodology – call it M – which includes other conceptual anchors besides pursuitworthy benefits and, hence, sometimes offers explications that differ from those of Pragmatic Conceptual Analysis. By definition, Pragmatic Conceptual

Analysis yields explications that will do optimally well to preserve the regular ways in which analysed concepts have delivered pursuitworthy benefits. Since M sometimes offers other explications, M must sometimes call upon us to sacrifice regular ways of achieving pursuitworthy benefits to preserve whatever features M takes to determine concept meaning (intuitions, metaphors or whatever). But when push comes to shove like this, Jackson's advice clearly applies: it is 'only sensible' to employ Pragmatic Conceptual Analysis rather than a competing methodology that would ask us to sacrifice pursuitworthy benefits in order to achieve something less worthy of pursuit.¹⁹

References

- Appiah, A. (1986) 'Truth Conditions: A Causal Theory'. In *Language, Mind and Logic*, Thyssen Seminar Volume, Jeremy Butterfield (ed.), Cambridge: Cambridge University Press, 25–45.
- Blackburn, S. (2005) 'Success Semantics'. In *Ramsey's Legacy*, H. Lillehammer and D. H. Mellor (eds.), New York: Oxford University Press.
- Boyd, R. (1988) 'How to be a Moral Realist'. In *Essays on Moral Realism*, Sayre McCord (ed.), Cambridge University Press, 181–228.
- Burge, T. (1979) 'Individualism and the Mental'. In *Studies in Metaphysics*, P. French, T. Uehling and H. Wettstein (eds.), Minneapolis: University of Minnesota Press.
- Capellen, H. (2012) *Philosophy without Intuitions*. Oxford: Oxford University Press.
- Carnap, R. (1950) *Logical Foundations of Probability*. Chicago: University of Chicago Press.
- Chalmers, D. (2002) 'The Components of Content'. In *Philosophy of Mind: Classical and Contemporary Readings*, David J. Chalmers (ed.), New York: Oxford University Press, 608–633.
- Chalmers, D. (2004) 'The Foundations of Two-Dimensional Semantics'. In *Two-Dimensional Semantics: Foundations and Applications*, M. Garcia-Caprintero and J. Macia (eds.), Oxford: Oxford University Press.
- Craig, E. (1990) *Knowledge and the State of Nature*. Oxford: Oxford University Press.
- Dennett, D. (1991) *Consciousness Explained*. Boston: Little, Brown.

¹⁹ This chapter overlaps heavily with ch. 2 of my book manuscript (Fisher n.d. 2), and both draw upon work in my doctoral dissertation (Fisher 2006). I am indebted to more people than I can name for helpful comments but especially Terry Horgan, David Chalmers, Joseph Tolliver, Chris Maloney and my colleagues at SMU, as well as audiences at the AAPA, the eastern APA, the SPP, the SSPP, the WCPA, many other combinations of letters, Florida International University and the Universities of Arizona, British Columbia and St Andrews. I am also grateful for helpful comments from other contributors to this *New Waves* volume.

- Dretske, F. (1988) *Explaining Behavior*. Cambridge, MA: MIT Press.
- Evans, G. (1973) 'The Causal Theory of Names'. In *The Philosophy of Language*, A. P. Martinich (ed.) New York: Oxford University Press, 1996.
- Fisher, J. (2006) 'Pragmatic Conceptual Analysis'. Ph.D. diss., University of Arizona.
- Fisher, J. (forthcoming) 'Pragmatic Experimental Philosophy'. *Philosophical Psychology*.
- Fisher, J. (n.d. 1) 'The Authority of Pragmatic Conceptual Analysis.' In preparation.
- Fisher, J. (n.d. 2) 'Pragmatic Conceptual Analysis'. Under review.
- Fodor, J. (1990) *A Theory of Content and Other Essays*. Cambridge, MA: MIT Press/Bradford.
- Frege, G. (1892) 'On Sense and Nominatum'. In *The Philosophy of Language*, 3rd edn, A. P. Martinich (ed.) 186–98. Oxford: Oxford University Press, 1996.
- Griffiths, P. (1997) *What Emotions Really Are: The Problem of Psychological Categories*. Chicago: University of Chicago Press.
- Haslanger, S. (1999) 'What Knowledge Is and What It Ought to Be: Feminist Values and Normative Epistemology'. *Philosophical Perspectives*, 13, 459–480.
- Hilbert, D. (1987) *Color and Color Perception*. Stanford, CA: CSLI Publications.
- James, W. (1906) 'What Pragmatism Means'.
www.marxists.org/reference/subject/philosophy/works/us/james.htm.
- Kornblith, H. (2002) *Knowledge and Its Place in Nature*. Oxford: Oxford University Press.
- Kripke, S. (1972) *Naming and Necessity*. Cambridge, MA: Harvard University Press.
- Lewis, D. (1983) 'New Work for a Theory of Universals'. *Australasian Journal of Philosophy* 61: 343-377.
- Lewis, D. (1984) 'Putnam's Paradox'. *Australasian Journal of Philosophy*, 62, 221–236.
- Margolis, E., and S. Laurence (eds) (1999) *Concepts: Core Readings*. Cambridge, MA: MIT Press.
- Millikan, R. (1984) *Language, Thought, and Other Biological Categories*. Cambridge, MA: MIT Press.

- Millikan, R. (1998) 'A Common Structure for Concepts of Individuals, Stuffs, and Basic Kinds: More Mama, More Milk and More Mouse'. *Behavioral and Brain Sciences*, 22(1), 55–65.
- Millikan, R. (2000) *On Clear and Confused Ideas*. Cambridge: Cambridge University Press.
- Nahmias, E., S. Morris, T. Nadelhoffer and J. Turner (2005) 'Surveying Freedom: Folk Intuitions about Free Will and Moral Responsibility'. *Philosophical Psychology*, 18, 561–584.
- Nahmias, E., S. Morris, T. Nadelhoffer and J. Turner (2006) 'Is Incompatibilism Intuitive?'. *Philosophy and Phenomenological Research*, 73, 28–53.
- Papineau, D. (1987) *Reality and Representation*. Oxford: Blackwell.
- Putnam, H. (1973) 'Meaning and Reference'. In *The Philosophy of Language*, A. P. Martinich (ed.), New York: Oxford University Press, 1996.
- Quine, W. V. O. (1969) 'Epistemology Naturalized'. In *Ontological Relativity and Other Essays*. New York: Columbia University Press.
- Ramsey, F. (1927) 'Facts and Propositions'. In *The Foundations of Mathematics, and Other Logical Essays*, R. B. Braithwaite (ed.), London: Routledge and Kegan Paul, 1931, 138–155.
- Rupert, R. (1999) 'The Best Test Theory of Extension: First Principle(s)'. *Mind & Language*, 14, 321–355.
- Stanley, J. (2005). *Knowledge and Practical Interests*. Oxford: Oxford University Press.
- Strawson, P. F. (1950). 'On Referring'. In *The Philosophy of Language*, 3rd edn, A. P. Martinich (ed.), Oxford: Oxford University Press, 1996, 215–230.
- Woodward, J. (2004) *Making Things Happen: A Theory of Causal Explanation*. New York: Oxford University Press.
- Whyte, J. (1990) 'Success Semantics'. *Analysis*, 50, 149–157.