The University of
Nottingham

UNITED KINGDOM · CHINA · MALAYSIA

McKeown-Green, Jonathan and Pettigrove, Glen and Webster, Aness (2015) Conjuring ethics from words. Noûs, 49 (1). pp. 71-93. ISSN 0029-4624

Access from the University of Nottingham repository:
http://eprints.nottingham.ac.uk/31991/1/Conjuring%20Ethics%20from%20Words%20%20%28Final%29.pdf

Copyright and reuse:

A note on versions:

Conjuring Ethics from Words

Sonya, on learning that Jasper is in hospital, wonders whether she should bring him flowers. Geoffrey the young architect wonders how he can make a name for himself and hypothesizes that if he wants to do that, he should straighten the Leaning Tower. Millicent the metaethicist wonders whether Sonya and Geoffrey are making judgments about matters of fact and, if so, whether these *normative* facts are anything over and above facts about the way the world is. She also wonders whether their judgments are the kinds of things that count as beliefs and whether the conditional or instrumental nature of Geoffrey's makes it fundamentally different from Sonya's. In the back of her mind she worries about goodness and rightness and how they enter into the analysis of these judgments.

Many of Millicent's colleagues address questions like hers principally by exploring the semantics of sentences like

1) Sonya should bring Jasper flowers; and
2) If Geoffrey wants to make a name for himself, he should straighten the Leaning Tower.

Indeed, this approach is common enough that Ralph Wedgwood can claim, "One of the central questions of metaethics is semantical. An answer to this semantical question would be an account of the meaning of moral statements" (Wedgwood 2001, 1; similarly Schroeter and Schroeter 2009, 1). Such a claim is unsurprising, given the history of metaethics and of the Linguistic Turn in 20th century philosophy more generally. The rise of expressivist and prescriptivist theories in the 40s and 50s, ongoing discussions of Moore's "open question" argument, and the like, lend credence to the claim that one of the central debates in the field has concerned semantics.

However, we shall argue that a number of metaethicists are asking semantics to deliver more than can reasonably be expected of it. For example, from an account of the semantics of 'ought', Stephen Finlay conjures conclusions about the teleological structure of moral reasons (Finlay 2009, 336). From claims about the semantics of English sentences, Frank Jackson derives the conclusion that ethical facts are reducible – necessarily and in principle *a priori* – to descriptive facts (Jackson 1998, 113ff.). From observations about "the semantics of 'ought' and 'good'", Wedgwood draws the conclusion that neither goodness nor rightness is metaphysically or explanatorily prior to the other (Wedgwood 2009, 517).

This sort of *direct* move from premises about the semantics of 'good', 'right', and 'ought' to conceptual or metaphysical conclusions about goodness, rightness, reasons, and obligations is puzzling. Linguistically-inclined metaethicists seldom offer a methodological justification for moving from linguistic premises to substantive conclusions and it is not clear that they would all offer the same justification, if pressed. We shall canvass the most promising justifications for making such a move. We show that these justifications fail and argue against current practice (in metaethics and elsewhere), which confuses an investigation of a word's meaning, reference, or competence conditions with an analysis of some concept or property associated with that word.

1. Locating the Problem

It is important to be clear, at the outset, about what is and what is not the problem we mean to diagnose. We do not deny that findings about the meanings, reference, or mastery of linguistic expressions can provide a useful resource for metaphysical inquiry. For instance, because philosophical investigation is characteristically pursued verbally, without the benefit of experimental manipulations of the non-verbal world, attention to the semantic details of verbal

formulations can be illuminating. In both professional and lay discussions of a philosophically fertile topic, entrenched reliance on a standard verbal formulation of some key thesis or question could lead one to overlook something loose or misleading in that formulation. Careful attention to features of the linguistic expressions that constitute the formulation might resolve ambiguities, reveal implicatures, or open up new lines of research into the topic itself.[1] Thus, we are not denying that there is a close relationship between semantics and metaphysics.[2] It is quite reasonable for metaethicists and other philosophers to investigate the semantic values, competence conditions and reference determination of linguistic expressions. Such semantic theorising can yield crucial insights into metaphysical issues, just as evidence from anthropology, physics and psychology can.

The metaethical projects that trouble us are not presented as projects in which one subdiscipline comes to another's aid. A telling illustration of the kinds of problems that can arise if one fails to keep a firm fix on the distinction between semantics, on the one hand, and a number of other things in which metaethicists might be interested – such as moral judgments, moral concepts, moral reasoning, and moral metaphysics – is provided by François Schroeter (2006). Schroeter's aim is to offer a critique of sentimentalism, which he characterizes as an "attempt to provide a credible account of the meaning of our most important evaluative terms .... I'll call 'sophisticated sentimentalism' ... the family of proposals which take the rational endorsement of our emotional responses to be the key to explaining the meaning of evaluative terms" (Schroeter 2006, 337 and 339). Already, this way of setting up the project should strike one as odd. It is anachronistic, to say the least, to suggest that the "traditional sentimentalists" – Shaftesbury, Hutcheson, Hume, and Smith – were concerned with "the main linguistic function of evaluative terms" (337; see Wiggins 1987, 187-188). But even as a characterization of what he calls "sophisticated sentimentalism" it is misleading. It may fittingly describe one of Allan Gibbard's and Simon Blackburn's metaethical interests.[3] However David Wiggins, in the piece Schroeter cites (Wiggins 1987), quite explicitly sets aside the discussion of words and their meanings – on which he thought emotivists and expressivists had misleadingly focused – in order to redirect attention toward moral thoughts, properties and judgments. Likewise, John McDowell's concern in the piece Schroeter cites is with "our

---

[1] For example, one strand of the debate about agency and moral responsibility considers a puzzle that tends to be formulated using a locution like 'the ability to do otherwise'. Hence, we find arguments like the following: "If my actions are causally necessitated, I cannot (it is not possible for me to) do otherwise. Yet, to have a free will, there must be some cases in which I can (in which it is possible for me to) do otherwise. So it would seem that I have no free will, given causal necessitation." Bruce Aune (1967) argued that the words 'can' and 'possible' are ambiguous and that I can truly say, given a certain disambiguation, that I can do otherwise, so long as I would have done otherwise had I chosen to. Aune's claim, and the debate it sparked with Keith Lehrer (1968), illustrate how an argument for the inconsistency of free will and causal necessitation could be influenced by reflection on a related semantic question. Similarly, a debate regarding hypothetical imperatives has focused on an alleged ambiguity in the locution, 'you ought, if you have the end, to take the means'. The discussion has centered on whether 'ought' takes wide scope (over the whole conditional) or narrow scope (merely over the consequent). Recently, Mark Schroeder has cast doubt upon the fruitfulness of this discussion, since it assumes, probably wrongly, that the relation expressed by 'ought' takes propositions for one of its relata (2004, 342-344). Here again, the investigation of an (allegedly) ambiguous linguistic expression is meant to shed light on a substantial, nonlinguistic, philosophical issue.

[2] To give another example, one might take a philosophical interest in sentences like: "Cannibalism is always bad," "The Leaning Tower is ugly," "You ought to straighten the leaning tower," and "The Leaning Tower is amusing." Apparently they have the same syntactic structure as fact-stating sentences that canonically express beliefs, such as, "Wood floats in milk," and "I can count to seven." This suggests that sentences in the first group are also fact-stating and that they too express beliefs. That in turn suggests that there are moral and aesthetic facts and beliefs. Exploring the semantics of these sentences *in conjunction with* a study of moral metaphysics might serve to reinforce or undermine this initial hypothesis. (See, for example, Sturgeon 2007).

[3] Even with them the fit is inexact. It does a better job of describing Blackburn's project in *Spreading the Word* than it does of characterizing the aim of *Ruling Passions*. In Gibbard's case, it would be a better gloss on *Thinking How to Live* than on *Wise Choices, Apt Feelings*.

equipment for thinking ethically" which he unpacks in terms of "our ethical concepts" (McDowell 1987, 156). Although it is clear that Wiggins and McDowell do not take their project to be the elucidation of "the main linguistic function of evaluative terms," they do occasionally make use of words or phrases that might also be at home in semantics, so it is perhaps understandable that someone who approached their work with an overarching interest in semantics might read them as sharing this interest. However, it is hard to see how Schroeter could construe the theory of value developed by Elizabeth Anderson (1993) as aiming to explain "the meaning of evaluative terms." She is concerned with respecting, appreciating, loving, admiring, and the like, and their import for reasoning and acting, not with 'respecting', 'appreciating', 'loving', and 'admiring', and their import for 'reasoning' and 'acting'.[4]

The issue is deeper than simply misdescribing the position he is criticizing. Having converted a number of quite different projects in moral epistemology, moral metaphysics, moral reasoning, and moral perception into a unified project in moral semantics, Schroeter periodically slides back into the language characteristic of one of these other projects, further obscuring any differences there might be between them. Thus, he shifts from "determining what falls into the extension of 'is morally wrong'" (349) to determining what "play[s] a decisive role in our reasoning" (350). And from observations regarding "what is necessary in order to count as 'conversable' with the word[s]" commonly employed in moral discourse (356), he draws conclusions about what do or "do not figure as decisive considerations in our most important evaluative judgments" (360). At the very least, before one is entitled to draw conclusions of the latter sort, one must explain why linguistic competence can be treated as a proxy for competence in evaluative reasoning.

Stephen Finlay (2009) provides an even more pointed example of the tendency to move from premises about linguistic meaning and semantic theory to conclusions about the nature of normativity. He begins from the thesis, "All instrumental ought-propositions can be transparently semantically expressed with an 'in order that ...' modifier" (2009, 318). (Hence, we might say 'In order that he make a name for himself, Geoffrey ought to straighten the Leaning Tower.') 'Ought' in such a sentence is explicitly relativised to an end. He then proceeds to argue that all occurrences of 'ought' can be similarly construed, so that every 'ought'-sentence has an implicit 'in order that ...' completion. Such completions are thus available for predictive sentences like: "He ought to be in Pisa by now" and apparently categorical sentences like: "We ought not to eat one another." One of the principal advantages of this account over others, Finlay suggests, is that it offers "a single reductive semantic analysis" for "predictive, instrumental, and categorical uses of 'ought'" (2009, 335). It is important to note that this relational theory is not intended to be merely linguistic.[5] It is a substantive normative theory that is both naturalistic and teleological. All obligations are relativized to some end or other. And normativity itself, he claims, is ultimately reducible to comparative probability relations between possible courses of action and a given outcome. This is not achieved by identifying a potentially misleading ambiguity in our moral language. Nor is it just an instance of noticing functional or grammatical similarities between different bits of a natural language. This is a case where considerations regarding what would make one semantic theory more parsimonious than another are taken as reasons for thinking one account of moral

---

[4] Even when she is talking about "meanings", as on page 3, it is clear that Anderson's concern is not with linguistic items. She is concerned with attitudes, reasons, valuations, and self-understandings. These are the things whose meanings are being interpreted. Of course, the language an individual or community employs when interpreting attitudes, etc., will have an influence on the resulting interpretation. So someone engaged in a project like Anderson's might also develop an interest in evaluative terms in their own right. But such a development would involve an explicit shift in focus to a different topic from the one with which Anderson is engaged.

[5] See note 11 below. Interestingly, concepts are introduced in Finlay's second thesis in a way that suggests they are indistinguishable from terms – or if they are not indistinguishable, it is not important to distinguish them (2009, 320).

metaphysics preferable to another. Surely this sort of inference stands in need of justification.

A third example of the kind of unmotivated shift with which we are concerned, from semantic premises to substantive moral conclusions, comes from Ralph Wedgwood (2009). The key he uses to unlock the mysteries of the priority relation between the right and the good is introduced when he turns his attention to "The semantics of 'ought' and 'good'" (2009, 507). From observations about ways in which meanings commonly associated with normative words like 'ought', 'right', and 'should' can be expressed using evaluative words like 'good' and 'better' (and vice versa), he draws the conclusion that neither goodness nor rightness is metaphysically or explanatorily prior to the other (2009, 517). Again, this transition stands in need of further justification. Why is one entitled to draw the conclusion that the good is neither conceptually nor explanatorily prior to the right (and vice versa) from the fact that sentences containing 'ought' can be transformed into sentences containing 'better' (and vice versa)?

2. Conceptual Analysis is Semantic Analysis
The authors we have discussed – and numerous others – write as though discovering the semantic values of linguistic expressions yields substantive conclusions about normative concepts or moral metaphysics. Why might someone think this? Leaving aside for the moment the more challenging task of deriving metaphysical conclusions from semantic premises, let us begin with the more modest supposition that there is a tight enough link between the semantic and the conceptual that one can draw conclusions about the latter from premises about the former.[6] One might think conceptual analysis *just is* meaning analysis and meaning analysis *just is* natural language semantics.[7] If this were one's starting assumption, then it might seem plausible to think that an investigation of the semantics of 'right' could yield an analysis of the concept of rightness.

When stated this baldly, it is not clear how many people would be inclined to endorse the proposed link between conceptual analysis and semantic analysis. First, the semantic analysis would need to focus on more than just a single word and the sentences in which that word appears. Consider, for example, the concept of a person. We can supply the analyst with linguistic data about our concept of a person without using a word that designates persons. We can do this by saying, "Those war crimes were committed sixty years ago, so punishment for

---

[6] How much more modest the latter supposition is will depend in part on what one takes concepts to be and how one construes the relationship between the conceptual and the metaphysical. If one takes a concept to be (a constituent of) a mental representation, then the gap between the conceptual and the metaphysical will be as wide or as narrow as the gap between the representation and what it represents. If one takes a concept to be a set of discriminatory and inferential abilities, then the gap between the conceptual and the metaphysical will be as wide or as narrow as the gap between such abilities and that on which they are exercised (whatever that might be). If one takes a concept to be something like a Fregean sense, then the gap between the conceptual and the metaphysical will be as wide or as narrow as the gap between these senses and their referents. (See Margolis and Laurence, 2011) In what follows we intend to be as ecumenical as possible regarding these and other competing accounts of the conceptual and the metaphysical. But as a general rule, the closer one takes the relationship to be between the conceptual and the metaphysical, the wider the gap will be between the semantic and the conceptual; and the closer one takes the relationship to be between the semantic and the conceptual, the wider the gap will be between the conceptual and the metaphysical.

[7] Frank Jackson defends this view. "Our subject is really the elucidation of the possible situations covered by the *words* we use to ask our questions" (Jackson 1998, 33; Jackson's emphasis). Jackson's contributions to metaphysics are explicitly treated as contributions to reference determination and competence conditions for terms in natural languages. However, he clearly sees his project as continuous with what is often called "conceptual analysis" and he adopts this label himself, partly in deference to the tradition. He also says that the label helps to emphasize that he is exploring "the elucidation of the various situations covered by bits of language according to one or another language user, or by the folk in general," and "is divorced from considerations local to any particular language" (1998, 33). Despite this concession to the traditional terminology, Jackson clearly thinks that there is no work for conceptual analysis to do that natural language semantics cannot already do unaided.

them is inappropriate now." Such linguistic data can supplement data gleaned from utterances containing the word 'person' as evidence about the situations covered by a certain concept – a concept that might well be associated with that word, but that can be signalled linguistically in other ways too.

Second, even if we broadened our linguistic analysis to include sentences in which the concept-word does not appear but that are relevant to its content, it is clear that we would not yet have attended to all of the pertinent data. The analysis of concepts is sensitive to nonlinguistic data as well.[8] Among the facts that settle which items count as persons are facts about which items we count in a census. Among the facts about which items you regard as red are facts about which traffic lights you stop for. The conceptual analyst will be attentive not only to what people say about persons and redness but also to other indicators of their conceptual commitments.

Third, the equation of conceptual analysis with linguistic analysis would overlook cases in which a person might grasp a concept even though she does not have a word or phrase with which to express it. For example, the concept of a person, of a cause, or of the emotion *amae* could be possessed by beings who had no word or phrase for expressing the concept. The concept could play a role in the thoughts, actions, and interactions of such beings. Such beings could, perhaps, distinguish people from non-people and approach problem-solving by positing causal processes. If our words for these concepts had not existed, these might have been concepts for which there are no words.[9]

For reasons like these, the first generation of 'ordinary language' philosophers were careful not to confuse their project with a merely linguistic one. John Austin in "A Plea for Excuses" offered the following cautionary note, "In view of the prevelence of the slogan 'ordinary language', and of such names as 'linguistic' or 'analytic' philosophy or 'the analysis of language', one thing needs specially emphasizing to counter misunderstandings. When we examine what we should say when, what words we should use in what situations, we are looking again not *merely* at words (or 'meanings', whatever they may be) but also at the realities we use the words to talk about: we are using a sharpened awareness of words to sharpen our perception of, though not as the final arbiter of, the phenomena. For this reason I think it might be better to use, for this way of doing philosophy, some less misleading name than those given above."[10]

3. Semantic Analysis Yields Conceptual Analysis

Although, for the reasons just mentioned, few will be inclined to *equate* conceptual analysis with natural language semantics, a number of people might endorse a slightly modified version of the link proposed above. One might think an investigation of the semantic properties of linguistic expressions *produces* among other things an analysis of the associated concepts. Hence, one might insist, we can accomplish conceptual analysis by doing linguistic analysis.[11]

---

[8] Joshua Gert (2011) overlooks this point despite showing a fleeting sensitivity to it: "One promising strategy for dealing with these threats to naturalism [from normative, mathematical, and other seemingly mysterious domains] is to move to a more sophisticated understanding of 'fitting into' the naturalistic world. This strategy begins by appealing to the kinds of facts that do not seem problematic ... facts about *human behaviour* and about our linguistic capacities .... This is the strategy I will pursue in the present paper, and I will call it 'linguistic naturalism'" (36; our emphasis). As his label would lead one to expect, he ignores non-linguistic behaviour in his subsequent analysis.

[9] This is consistent with the possibility that we acquire many concepts by learning words that express them and with the possibility that there are some concepts that could only be acquired in this way.

[10] Austin, "A Plea for Excuses," *Philosophical Papers*, 3rd ed. (Oxford, 1979) 182.

[11] This appears to be Wedgwood's position: "As a philosopher, I am primarily interested, not in empirical questions about the meanings of words, but in the nature of the *concepts* that those words can be used to express – especially when those concepts are central to certain branches of philosophy, as the concepts expressed by 'ought' are central to ethics and to the theory of rational choice and rational belief. Still, it is often easiest to

However, we shall argue that even when a concept is semantically associated with a particular linguistic expression, we should not assume that, therefore, the nature of the concept is fully revealed by linguistically represented facts about the expression. There are two points, in particular, that we want to emphasize. First, although a semantic analysis of an expression may at times yield an analysis of the concept(s) associated with it, sometimes the causal relation runs in the other direction. Sometimes – perhaps often – we need to have mastered the concept first before we are in a position to provide the semantic analysis. Second, there is a limit to the conceptual information that can be obtained by pulling expressions apart. As a result, there may be conceptual complexities that are not readily disclosed in the linguistic structure.

Let us start with a simple physical case before we turn to more controversial normative ones. English expressions like 'the local drugstore' are syntactically well-formed. Standard semantic theory tells us that their semantic properties are therefore generated in a rule-governed way from the semantic properties of their components – more precisely, the components of their logical forms – and, plausibly, these include the concepts associated with 'local' and 'drugstore'.[12] So the semantic structures of those expressions arguably reveal some internal structure of that concept. For the same reason, competence with the concept LOCAL DRUGSTORE might be explicable in terms of semantic competence with 'local', 'drugstore' and the relevant, linguistically revealed, rules of composition.

Suppose, however, that I am interested in analysing what it is for something to be, not a local drugstore, but a drugstore. Perhaps I am a New Zealander learning an American dialect. Although some American expressions translate New Zealand ones (e.g. 'soda' in the Northeast, 'coke' in the South, and 'pop' in parts of the Midwest apply to what New Zealanders call a 'fizzy drink'), other American words, like 'drugstore', apply to things that are unfamiliar to New Zealanders. I might attempt to advance my understanding of 'drugstore' by pulling the word apart into the constituents 'drug' and 'store'. However, this will not get me very far. It may enable me to distinguish between a drugstore and a drug manufacturer. But one will still be a long way from mastering either the concept DRUGSTORE or the lexeme 'drugstore'.

A better plan will be to look at sentences in which the word appears and to see how they might be translated into New Zealand English. Some of the sentences in which 'drugstore' is used are adequately translated by those in which a New Zealander might use the word 'chemist': for instance, 'The doctor told me to go to the drugstore to fill a prescription'. In other sentences, however, 'chemist' would clearly not be an adequate substitute and in those sentences the New Zealand 'dairy' seems a more suitable replacement: for instance, 'I'm going

---

approach the task of giving an account of the nature of certain concepts by studying the meanings of the words that can express those concepts." (2006, 127). It becomes clear as one moves through the paper that Wedgwood does not merely think one can approach concepts via semantics but that the latter generates the former. Once one has analysed the words it appears there is nothing else for the conceptual analyst to do. Finlay adopts a similar position, although in his case what is generated are metaphysical conclusions. In "What *Ought* Probably Means, and Why You Can't Detach It," (*Synthese* 177 [2010]: 67-89) he criticizes "the standard semantics" of 'ought' on the grounds that, "We might have hoped that discovering the meaning of 'ought' would teach us something about the mysterious nature of normativity itself," but the standard semantics of 'ought' leaves normativity mysterious (77). His own semantic analysis of 'ought', by contrast, does tell us something about normativity. Indeed, he suggests that the version of the analysis offered in "Oughts and Ends" (2009) aims "to establish the truth of" his theory of normativity (68).

[12] The logical form of a sentence is the level of syntactic representation that reveals the meaningful components of the sentence and the semantic relations that hold among them. In generative grammar, it is contrasted with other levels of syntactic organisation, such as phonetic form, which reveal different (non-semantic) aspects of the structure of a sentence. Many 'surface level' constituents of sentences are re-ordered, disassembled, or dispensed with at this semantic level of description. Logical forms often contain unrealised constituents, including tense markers and traces. Semantic values (ordinary objects, truth values, functions and so on) are assigned only to constituents of logical forms. "Whatever the nature of LF, it is supposed that all grammatically determined information that is relevant to interpretation is to be found there" (Higginbotham 1985, 549).

to stop off at the drugstore to pick up some milk and a candy [chocolate] bar.' In yet other sentences, neither 'chemist' nor 'dairy' will work as a substitute and, in those sentences, 'drugstore' is better translated by 'bottle shop': for instance, 'I'm going to run to the drugstore to pick up some beer.' So 'drugstore' is a bit like 'chemist,' a bit like 'dairy,' and a bit like 'bottle shop,' but even combining these terms with each other and with others that might suggest themselves – 'chemist cum dairy cum bottle shop cum stationer cum hardware shop cum clothing shop' – fails to capture satisfactorily the semantic properties of 'drugstore', the conceptual properties of DRUGSTORE or the metaphysical properties of drugstores.

When a New Zealander becomes competent with 'soda', she gains semantic but not conceptual knowledge. By the time she has mastered 'drugstore', she will have gained conceptual as well as semantic knowledge. However, it is misleading to suggest that the semantic project will have yielded the conceptual result. The causal process is likely to have gone in the other direction. She is only likely to be in a position to identify the semantic properties of 'drugstore' after – and partly as a result of – visiting one or more American drugstores and speculating about the classificatory or historical facts that explain which goods are sold there. The successful completion of the semantic project will depend upon the analyst's prior conceptual mastery.

Turning now to a normative example, in a series of recent articles Stephen Finlay has argued for a comparative probabilistic, end-relational "theory of the semantic core of 'ought'" that he thinks can "teach us something about the mysterious nature of normativity itself."[13] What he thinks it reveals is, among other things, that normativity is itself end-relational. That is, it has a teleological structure that, he suggests, can accommodate both Utilitarians and Kantians: "Utilitarianism, for example, is animated by the perception that what we morally ought to do can be identified with what is most promotive of maximal human wellbeing impartially considered, and Kantianism, with what we must do in order that we are able to will our maxims as universal law."[14] However, this substantive claim about normativity (or, even more modestly, about our concepts of normativity) does not just fall out of the semantic analysis. Michael Stocker (1969) observed several decades ago that deontological principles can be recast in a teleological, consequentialist form. This translation exercise may tell us something about the elasticity of language. It may even tell us something about the kinds of thoughts we entertain when deliberating about what to do. But it does not yet tell us about normativity, for it does not yet tell us whether anything has been lost in the translation. Stocker has argued at length that something is in fact lost: it obscures the normatively significant distinction between that *from* which we act and that *for* which we act.[15] As it happens, we agree with Stocker, but that is not the point. Rather, the point is that in order to construct his semantic analysis of 'ought' Finlay has to decide whether or not the way in which he has recast the instrumental ought, the predictive ought, the categorical ought, the subjective ought, the objective ought, etc., is adequate or whether something has been lost or altered in the reformulation. And in order to perform this assessment, he has to draw on prior conceptual commitments. The conceptual conclusions *yielded* by the semantic analysis are the conceptual commitments that were *presupposed* by the semantic analyst. Someone who rejects those conceptual commitments is also likely to reject the resulting semantic analysis.

A second normative example, this time relating to the limit on the conceptual information that can be obtained by pulling expressions apart, is offered by the much debated relationship between the word 'flaw' and the concept of a flaw.[16] A student of the rules that govern the relationships between English expressions and the world will eventually face the metaphysical

---

[13] Finlay (2010) 68 and 77.
[14] Finlay (2009) 336.
[15] See especially Stocker (1981).
[16] For crucial early contributions to this debate, see Chomsky (1981) and Higginbotham (1985).

question of what, if anything, a flaw is, just as she must face the similar question about what a drugstore is. Facts about the nature of drugstores may be fairly uncontroversial, but this is not true of flaws. Somebody who believes that the referent of a lexeme had better be a physical object might prefer that 'flaw' not be a lexeme. Any such problem can be avoided if the truth of a sentence like, 'There is a flaw in Chomsky's argument' did not require that flaws exist. It will therefore be convenient if the logical form of that sentence is something like: 'Chomsky's argument is flawed', which replaces the troublesome noun phrase with a predicate. Plausibly, though, this fix cannot be extended to cover sentences like: 'We managed to deal with three of the flaws in the argument, but the other two are still causing trouble.' Reformulations are available that extend the predicative analysis to this sentence, but fresh counterexamples loom. One wonders whether each new repair can be justified as part of a unified theory. As Peter Ludlow says, "There is a heavy burden to show that the introduction of [any] proposed hidden structure comports well with the rest of what we know about [the] syntactic form [of] these constructions" (Ludlow 2003, 149). One might be attracted to more revisionist proposals about the logical form of the original sentence, such as: 'Chomsky's argument does not measure up to the appropriate standard.' Deriving this paraphrase from the original sentence involves conceptual theorising about the role played by the word 'flaw' in sentences containing it and denying that it, or any associated predicate, is a minimal component of the logical form. But even if we go this way, "it has to be shown that general rules are available and not just case-by-case fixes" (Ludlow 2003, 149).

Scepticism about the possibility of doing this suggests a different path: perhaps our failure to find linguistic structure here suggests that there is none to find, just as there is none to find in the case of 'drugstore'. Ludlow, for instance, advocates settling the debate by adopting the semantic hypothesis that 'flaw' is a lexeme whose extension is the class of flaws. Such a solution leaves the conceptual analysis of the concept of a flaw, the question of what it takes to master that concept and the question of whether flaws are metaphysically possible to those who are more interested in flaws than in 'flaw'.[17] It lets us distinguish conceptual and metaphysical issues about the conditions something must satisfy to be a flaw, on the one hand, from linguistic issues about the rules that relate the word 'flaw' and its synonyms to the real or imagined things (including flaws) about which we can speak and write, on the other.[18]

Conceptual analysts who invoke linguistic theory to motivate metaphysical claims write as though the internal structures, if any, of concepts like GOODNESS, REDNESS, or WATER can, in principle, be revealed in the logical forms of sentences containing words like 'good', 'red', or 'water'. But if such words correspond to atomic components of logical forms in the way that, according to Ludlow, 'flaw' does, there is no mandate for treating the internal structures or mastery of philosophically interesting concepts as aspects of linguistic structure or mastery.[19]

To be sure, it is sometimes unclear or indeterminate whether some aspect of internal structure is linguistic or merely conceptual. Consider the word 'must' and pretend that its semantic contribution to containing sentences is accurately captured by standard interpretations

---

[17] Of course, there is no reason why a semanticist cannot pursue these. She may well wish to, since mastery of the concept is required for mastery of the word and an understanding of what flaws are completes the account of how reference is determined. Our point is merely that any structure that the concept may have - such as the structure that might distinguish flaws from hitches - is not represented linguistically, according to Ludlow's hypothesis.

[18] Of course, an artificial language that was designed to include metaphysical simples or unanalyzable concepts in its domain of quantification would presumably allow paraphrases of all the sentences containing 'drugstore' and all the sentences containing 'flaw' into sentences that revealed the metaphysical structures of drugstores and flaws or the structure of the associated concepts.

[19] Of course, we could assimilate all conceptual structure to linguistic structure if the parts of the mind involved in concept acquisition and comprehension also handle linguistic competence or if there is a tight logical connection between the study of language and the study of concepts. However, it is controversial to claim either of these things and the former possibility is surely an empirical claim awaiting more research.

of the box symbol from modal logic.[20] Very roughly speaking, one might want to say the following about the semantics of 'must': a sentence of the form 'It must be the case that p' is true *as a matter of semantics* iff the clause within the scope of the 'must' operator in the sentence's logical form (the clause represented by p) is true at every relevant (or accessible) world. This analysis treats the lexeme 'must' as semantically complex. Alternatively, one might say the following: 'It must be the case that p' is true *as a matter of semantics* iff p is necessarily true, in some special technical sense. One would add that the question of what it takes for p to be necessarily true in this sense is a metaphysical question, perhaps one on which conceptual analysis and standard modal systems can shed some light. This alternative analysis treats the lexeme 'must' as atomic with respect to what language can tell us about any concept it may express. Competing analyses of this kind may be available for many 'grammatical words'. We will not consider whether there is a fact of the matter about where linguistic facts yield to language-independent conceptual facts. We have argued, however, that there are non-grey areas and that some of the standard preoccupations of conceptual analysts may well fall within them.[21]

In sections 2 and 3 we have argued against treating semantic analysis as if it were conceptual analysis or as if it directly yielded conceptual or metaphysical conclusions. As we have noted:

a. the analysis of a concept may rely on linguistic data in which no word for that concept occurs,
b. conceptual analysis is sensitive to nonlinguistic data,
c. there are concepts that are not semantically associated with any particular word,
d. concepts might be accessible to beings that have no way to express them linguistically,
e. semantic analysis is often dependent upon prior conceptual knowledge, and
f. concepts may have structural features that are not represented as structural features of language.

These are challenges that would need to addressed before one would be justified in treating conceptual analysis as though it were either nothing but linguistic analysis or as if it were reliably generated by a semantic analysis. But, as we shall argue in the next three sections, there are additional obstacles that the semantically inclined metaethicist will also need to get around. And these may prove even more challenging.

4. Word Usage and Word Meaning

Perhaps the most challenging obstacle in the way of the semantically-inclined metaethicist emerges when one considers the semantic-pragmatic distinction.[22] If our interest is in the concepts that words can *be used* to express,[23] the analysis of the relevant linguistic data should,

---

[20] This pretence is for illustrative purposes only. The illustration could, we think, be adapted to the more sophisticated and more plausible semantics for modals developed by Angelika Kratzer (1981).

[21] This point, about how linguistic facts about the *content* of an expression may not exhaust the conceptual facts about concepts that might constitute that content, can be illustrated by an analogy. Alongside their interest in the content (or meaning) of expressions, semanticists are interested in the reference of expressions. It is thus the semanticist's business to ascertain how the reference of 'Kaplan' is determined and to discover that Kaplan, rather than some other object, is the referent of 'Kaplan'. Unless she is also a metaphysician, however, it is not her core business to explore the nature of Kaplan himself.

[22] To help fix ideas, here is a way to capture our usage of "semantic" and "pragmatic". "Semantic information is encoded in the sentence; pragmatic information is generated by, or at least made relevant by, the act of uttering the sentence." (Bach 2002, 284.) Crucially, information encoded in a sentence can, but need not, be among the information made relevant by an utterance of that sentence on a particular occasion. Thus, as we are construing it, the domain of the pragmatic has to do with all aspects of performing and interpreting communicative acts and some of those aspects draw on semantic information. Appeals to the pragmatic, in this sense, are thus not merely attempts to explain what is left over when semantics runs out.

[23] This wording is borrowed from Wedgwood (2006) 127.

in the first instance, be part of an analysis of *language use*, rather than of *language meaning*.[24] When we draw on what you say about some issue in order to learn what you think about it, we are interested primarily in *what you mean by what you say* rather than in *what your words mean*. If you use an unfamiliar word, we might consult a dictionary to find out *what it means*, or engage in some semantic theorising, but only so as to determine *what you meant by it* and we should probably not always assume that you used it literally. What one means does not always include what one's sentence means. In fact, it is not obvious even that it usually does.

All of this applies when we are gathering, introspectively or otherwise, linguistic data about conceptual commitments. Information about which actions the folk count as right, for instance, comes from reflection on what the folk say, or would say, about rightness *in various contexts*. The conceptual analyst cares about what the folk mean by what they say, or would say, in such contexts. Pragmatic facts about when and under what circumstances we say of an action that it is right shed light on the conditions under which we regard something as right and hence contribute, alongside facts about nonlinguistic activity, to a picture of our concept of rightness.

Of course, information about communicative and expressive acts is also deployed in the quest for *semantic facts.* This often requires controversial assumptions about which uses of words are literal. It may also require supplementing data about the acts we perform by uttering linguistic expressions on particular occasions with information about what our informants take the meanings of words to be, together with general principles of the kind popularised by Grice, which help to distinguish the the information semantically encoded in an expression from other information that is only made relevant by utterances of it. Such semantic theorising may be one significant tool in the service of the linguistically focused conceptual analyst's pragmatic investigation. It may sometimes be necessary to distinguish semantic from other contributions to language use in order to make sense of facts or intuitions about language use. However, the linguistically focused analyst should not regard the formulation and defence of semantic hypotheses as her central task; it is at best a means to a pragmatic end.[25] Furthermore, although facts about language use and our reflections on it inform both semantic theorising and conceptual analysis, there is no reason to assume that they inform the latter *only because* they inform the former.

Nevertheless, some semantically-inclined metaethicists might think that at least sometimes we are justified in holding that we can analyse concepts simply by doing semantics, because they think there are important cases where the pragmatic task reduces to a semantic task. This

---

[24] What we say here also responds to the following point made by Matthew Chrisman. "Some will want to switch to talk of concepts instead of words. That is, they'll suggest that metaethicists should give up on developing a semantics for normative words like `ought' and focus on something like the nature and content of the concept expressed by the all-things-considered normative `ought'. ... However, theorizing in both ethics and semantics seeks reflective equilibrium based in large part on input from ordinary intuition, and this input comes *couched in language using mostly ordinary words like 'ought'* [our emphasis] and not neologisms like 'all things considered normative "ought"'. So, even if metaethicists are most interested in the logic of our concepts rather than the meanings of our words, and these are projects kept at arm's length, I still think they should seek to integrate their account of the relevant concepts with a full understanding of the best linguistic theories of the meaning of the words used to express them" (forthcoming). As with Wedgwood, we think that Chrisman is drawing attention here to the significance of *language use*, rather than directly drawing attention to the *meanings of words*.

[25] Joshua Gert (2011) at times appears to attend to these sorts of pragmatic questions when he insists that we consider the conditions under which moral discourses might emerge in the first place and how subsequent generations might be inducted into them. (His interest in these questions is motivated, in part, by the desire to show that our moral discourses and related practices are unproblematic from a naturalist's point of view.) However, Gert approaches this pragmatic project by considering how individuals become semantically competent with particular expressions (such as 'morally wrong') and how these expressions gain their meanings and are "preserved from one generation of language-speakers to the next" (55). At the very best the semantic information on which he focuses would not tell us about how the discourses themselves emerge and are maintained because they leave out all of the pragmatic phenomena that shape what we manage to express and communicate with language. Hence, it fails to address the questions he sets out to answer.

happens, it might be urged, when linguistic expressions are sincerely uttered and intended literally. Such assertions tell us directly about the contents of the beliefs they express.[26] If I sincerely assert, "I know that he is lying," and intend this literally, then I believe that I know that he is lying and the content of the sentence is the content of the belief. Hence, if we can settle what it takes for an utterance of that sentence to constitute a sincere, literal assertion and if we have a theory of the semantics of this sentence, we have a theory about what the expressed belief commits one to. Less ambitiously, if we have a semantics for the sentence, we will know some interesting facts about the content of any belief that would be expressed by a sincere assertion, consisting of a literally-intended utterance of that sentence. Either way, we will have some information about our conceptual commitments regarding knowledge. Thus, in the case of this sentence and plausibly many others, it might be suggested, understanding the semantics of a linguistic expression is the central task in understanding conceptual commitments.

Let us grant that the content of a declarative sentence is identical with the content of the belief that is expressed when that sentence is asserted sincerely and intended literally. If this identity principle is to tell us anything substantial about the content of a belief, we must know what counts as a literal use of the relevant sentence. We need to know, for instance, which aspects of what a speaker means on a particular occasion are due to scalar implicatures, which are metaphorical, which are ironic and so on. This involves pragmatic, as well as semantic, considerations. Take the claim that, because of the semantics of 'know' and 'believe', knowledge entails belief. In a situation where I sincerely say, "I know that he is lying," and intend that sentence to be interpreted literally, I would not sincerely say, "I believe that he is lying," and mean it literally. This is because the latter utterance would be conversationally inappropriate.

In some cases sorting out the pragmatics of sincere and literal utterances may be relatively simple and uncontroversial. For example, the knowing/believing case can be explained by something like Grice's first maxim of quantity (Grice 1975/1989, 26). However, this does not mean that by assuming they are dealing with a sincere and literal assertion the linguistically-inclined metaethicist can safely ignore pragmatic issues and focus on semantics. There are three reasons for continuing to press this point. First, there are important cases where the pragmatic back-story for a semantic claim is both complicated and contentious. For instance, in order to defend the extension of his (indexical) end-relational account of the meaning of 'ought' to categorical ought-sentences like, "We ought not to eat one another," Stephen Finlay appeals to pragmatics.

> Here we encounter a *rhetorical* device. To speak in a way that supposes something to be true of your audience that is clearly controversial or false is a way of expressing a demand that it *be* true of them; consider 'In this family, we do not belch at the dinner table!' and 'You *will* come here!' By categorical use of 'ought', therefore, a speaker expresses the demand that his audience share his concern for the relevant end, and consequently for the behaviour at issue. ... Categorical uses of 'ought' are rhetorical uses of the end-relational 'ought' based on ellipsis (2009, 333).

Finlay's conjecture is meant to show that, despite appearances to the contrary, all apparently categorical utterances of 'ought'-sentences are end-directed. One's assessment of that conjecture will depend on which assertive uses of 'ought'-sentences one takes to be literal and sincere. And determining that will require one to do a fair amount of pragmatic, conceptual,

---

[26] Presumably, if the sentences concerned contain indexicals, anaphora, attributives, elliptical constructions, or any other expressions whose semantic import depends on features of context, those features need to be readily recoverable from the context of utterance. We ignore this complication here, but it is important, since some accounts of philosophically interesting terms, like 'ought' (Finlay, 2009), 'knowledge' (DeRose, 2005) and 'water' (Jackson, 1998) take important semantic features to be relativised to context.

and metaphysical work. Thus, the semantics of 'ought' sentences cannot straightforwardly be taken to deliver conceptual conclusions about reasons, obligations, or normativity.

There is a second problem for the argument from the possibility of sincere, literal assertion to the claim that pragmatic influences can be minimised in important cases so that the semantics of an expression gives us access to conceptual commitments. The problem is that there are many philosophically interesting conceptual projects that are not rescued by this argument. Sincere, literal assertions in which the word 'person' occurs are arguably not especially good guides to our concept of a person, since the folk do not reliably use this word in its philosophical sense. If Raz is right, the concept of law that is sought by philosophers of law is not captured by the semantics of 'law' or any other English word (Raz 2001, 7). If Matthew Chrisman is right, the deep questions regarding the nature of morality are not, and are not isomorphic with, questions about the semantics of natural language. He takes 'ought', like 'must' and 'may', to be an indexical modal auxiliary whose semantics do not settle standard metaethical debates. Rather, the issues that divide realists and expressivists, he argues, must be addressed by asking meta-semantic questions about why 'ought' and related words have the semantic values they do (Chrisman, forthcoming). The content of "There ought to be clowns", "Sonya ought to bring Jasper flowers," "Geoffrey ought to be in Pisa by now," or "We ought not to eat one another," may well be the same as the content of some belief that the sentence can literally and sincerely express. Nevertheless, if Chrisman is right, this belief may or may not reveal distinctively moral conceptual commitments. Whether it will depends on facts about the values of indexicals in particular semantic contexts and, even then, the revealed conceptual commitments will not settle the sorts of metaphysical questions that metaethicists consider.

There is a third reason for thinking that the assumption of a literal and sincere utterance will not enable the metaethicist to derive the kinds of conclusions she wants from a semantic analysis of some locution. It arises from the fact that even in uncontroversial cases all that these assumptions yield are facts about the belief expressed by the speaker (or some community of speakers). If one takes a concept to be (a constituent of) a mental representation, then this might yield a conceptual conclusion. But if one takes a concept to be a set of discriminatory and inferential abilities or if one takes it to be a Fregean sense, then a semantic analysis of a locution uttered under these conditions may not yield a conceptual conclusion. The semantic analysis might get us to our belief, but the concept might outstrip our belief. And the reality might outstrip both concept and belief.

Because of the three problems we have just discussed, the principle that the content of a sentence is the content of a certain sort of belief cannot rescue the claim that semantic analyses of the kind that characterise much current metaethical research play a central role in conceptual analysis. Sometimes, perhaps, the pragmatic facts of a particular situation are sufficiently straightforward to allow an analyst to recover crucial facts about the contents of mental states from a semantic analysis of particular expressions. But often they are not. Thus, at the very least, it will be incumbent on the metaethicist who looks to semantics for substantive normative insights to demonstrate that the case she is dealing with is one of the straightforward cases. And this will require her to assemble non-semantic evidence as well to support her conclusion. So at most the semantically inclined metaethicist can claim the following: linguistic data – both semantic and pragmatic – can be useful for conceptual analysis.

5. There are Paradigmatic Ways to Express Concepts

Consider another strategy for defending the view that investigating a concept is more or less the same as investigating the meaning-related properties of some linguistic expression. This strategy involves claiming that the use of certain expressions amounts to a canonical, standard, or paradigmatic way of expressing some concept. Consequently, exploring the semantics of such an expression is an effective means of analyzing the associated concept. For example, Matthew

Chrisman asserts, "Ethical theorists are interested in the meaning of the word 'ought' largely because the paradigmatic way in English to state general moral principles as well as specific practical conclusions is with an 'ought'-sentence" (Chrisman 2011, 1).

What is it for an expression to be a *paradigmatic* way to state something? If Sonya should bring Jasper flowers, we may say that she should, or that she ought to, or that she must, or that she had better, or that it would be best if she did; in the right circumstances, a change of facial expression will do it. Could the most frequently used be the one we are looking for? Evidently not. A particular expression could be reliably associated with a communicative purpose even when its semantics tells us nothing about that purpose. Consider, 'How are you?'

Perhaps to claim that 'ought'-sentences are paradigmatic is simply to claim that they semantically encode some concept that is central to our understanding of morality: when used literally, uttering an 'ought'-sentence amounts to talking directly about the concept, so that studying the semantics of 'ought' is much the same thing as studying that concept. The problem with this reading of "paradigmatic" is that, as we argued at the end of the previous section, it is sometimes not obvious that the words in which philosophers take an interest are paradigmatic, in this sense. It is only by doing some analysis, with the help of both semantics and pragmatics, that a robust hypothesis emerges about the semantic content of words like 'ought'. As we noted, Chrisman himself (forthcoming) argues that the concept expressed by 'ought' is perhaps not the normative one that some ethical theorists are arguing about, but rather a general purpose modal one. So if we adopt this 'literal use' gloss of 'paradigmatic', Chrisman himself is committed to denying that 'ought'-sentences are paradigmatic.

Is there a strategy for developing this idea into a justification for treating conceptual analysis as natural language semantics, that avoids the objections we have already discussed?[27] One strategy worth exploring, although it ultimately fails, involves accepting that the content of a natural language word is sometimes determined by the conceptual role played by that word in a folk theory. This is, in effect, a particular story about how to negotiate all the pragmatic phenomena involving the word, in order to arrive at an account of literal meaning. If the story were plausible, it would vindicate the idea that the ordinary words in which analytic philosophers take an interest – like 'belief', 'person', 'right' and 'should' – behave semantically as labels for precisely the concepts or categories that one might expect them to. It is in that sense that the uttering of sentences involving those words would be *paradigmatic*. We are not convinced by the story, but in order to explain why, we will first try to motivate it.

The first step towards endorsing a *folk-theoretic conceptual role semantics* is to accept the view that families of conceptual commitments can be treated as *folk theories*. Functionalism is sometimes presented as the application of this view to the metaphysics of the mind and is a useful illustration of the general strategy. For the functionalist, a mental state, like the belief that Grice is nice, is whatever (physical or other) state plays a certain characteristic role in deliberation, as well as in the production, prediction and explanation of behaviour. Introspection tells us what this role is, because introspection tells us what would count as a belief or a desire. Intuitions about networks of mental states, perceptual inputs, and behavioural outputs are described as comprising a folk theory which, in turn, is regarded as the kind of thing that could, in principle, be represented by a set of sentences closed under logical consequence. In that respect, it is apparently just like any formalisable scientific theory. I have an intuition that if you tell me Grice is nice, I will, other things being equal, come to believe that Grice is nice. The thought is that this intuition is a clause in my folk theory of the mind. This formal treatment is encouraged by the fact that I can express intuitions of this sort linguistically, by saying, for instance, that a particular possible case, described in a particular way, is one in which it would be fitting for me to say, "I believe that Grice is nice". This has

---

[27] We are not saddling Chrisman with the claim that there is such a strategy. He does not seem to be committed to the idea that linguistic premises can licence conceptual conclusions.

encouraged philosophers to treat English words like 'belief' and 'desire' as terms in our folk theory of the mind. So just as 'force' and 'momentum' receive their technical Newtonian meanings from structural features and interpretations of Newtonian theory, so 'belief' and 'desire' receive their semantic properties from their folk-theoretic roles. To find out about momentum *just is* to find out about the truth conditions and other semantic properties of equations and other Newtonian sentences containing the term 'momentum'. Similarly, on this construal of functionalism, to find out about desires *just is* to find out about the semantic properties of sentences containing the word 'desire', where those properties are revealed by interpreting the (natural) language in which the folk theory of the mind is expressed.[28] This *folk-theoretic* approach to intuitions with its characteristic construal of natural language semantics has analogs in other sub-disciplines.[29] Folk theories of truth, natural kinds, colours and morality have been proposed.[30] Jackson and Pettit (1995) describe their account of the latter as moral functionalism.[31] Lots of common nouns (like 'belief'), adjectives (like 'good') and auxiliaries (like 'ought') are thus taken to receive their semantic properties from their roles in clauses of folk theories, so that, in effect, they become labels for theoretical concepts. If this folk-theoretic account of their semantics is correct, we could intelligibly say that uttering sentences in which they occur is a paradigmatic way to assert clauses of the theories and this would vindicate the view that investigating the folk commitments theorised about is of a piece with investigating the semantics of the ordinary expressions used to talk about the phenomena.

We reject this argument for folk-theoretic semantics about words in a natural language. We are happy to grant that the folk's conceptual commitments regarding, say, the mental or moral might constitute a folk theory. However, the fact that the folk can express tenets of such a theory by *using certain words* is primarily a matter of pragmatics. Developing it into a view about sincere, literal utterances and hence, into a view about the semantics of those words involves the difficulties to which we alluded at the end of the previous section. Of course, a metaphysician who wants to understand and explore a folk theory could do with terminology to label its theoretical posits. Just as physicists attach particular meanings to 'spin' and 'momentum', so philosophers of mind or metaethicists will need words to pick out folk-theoretical concepts like those that the folk talk about. The defender of a folk-theoretic semantics for natural language words apparently assumes that the semantic properties of words that are routinely used by the folk to express folk-theoretic claims about certain concepts are the same as the semantic properties of the *theoretical terms* with which a metaphysician studying the folk theory might label those concepts. We see no reason for this assumption. The metaphysician might requisition the word 'belief' as a term of art – just as physical theory includes 'heat' among its terms of art – and stipulate that it labels whatever plays the appropriate functional role. This does not mean, though, that the content, reference, or competence conditions of 'belief' in English are given by this role which, after all, can be discussed by the folk using all sorts of locutions and is manifested by nonlinguistic behaviour too.

We allow, then, that semantic analysis of some artificial language, like a language designed to capture a folk theory, would be equivalent to conceptual analysis.[32] However, we deny that

---

[28] See Braddon-Mitchell and Jackson (2007).

[29] Although functionalism has been influential in the development of these analogs, connections between folk commitments and scientific theory, as well as the linguistic conception of a scientific theory, are much older. See, in particular, Carnap (1956).

[30] See Jackson (1998). See also Lynch (2005) for a functionalist theory of truth.

[31] For a different, but equally folk-theoretic, account of morality, see Smith (1994).

[32] Some regimented versions of (fragments of) a natural language might qualify. For instance, one might try to explain the concept of a rational requirement by giving a formal semantics for 'rational requirement' and regard this as a conceptual analysis of a notion that is implicit in folk practice: one that a philosopher might refer to by using the English phrase 'rational requirement' (even if the folk seldom do). One might think that such semantics

folk-theoretic approaches to the mental, the moral, and the like licence us to regard *ordinary natural language expressions* like 'belief' and 'ought' as theoretical terms in a folk theory.[33] So we are no closer to understanding what it means to say that "the paradigmatic way in English to state general moral principles as well as specific practical conclusions is with an 'ought'-sentence." More generally we have been unable to develop the appeal to a canonical, standard, or paradigmatic locution into a justification for the type and degree of attention that metaethicists have lavished on the semantics of 'ought', 'good', 'right', and 'rational'.

6. Externalism Closes the Gap

The final argument that we shall consider for thinking semantic analysis might directly yield substantive metaethical conclusions is derived from an externalist position regarding semantic properties that develops the insights of Kripke (1972), Putnam (1975), Donnellan (1970) and Kaplan (1989). An externalist of the sort we are imagining might claim that 'Geoffrey' refers *directly* to Geoffrey without any mediating description. Similary, they would claim that 'water' connotes H_2O, without this being due to some criterion that is semantically associated with 'water' and uniquely satisfied by H_2O. They would claim that the referents and extensions of these and other expressions are determined by causal relations that connect agents who have mastered the expressions with the relevant items.

Now we have argued that semantics is not, all on its own, a satisfactory way to do conceptual analysis. This apparently matters to philosophy at least partly because the exploration of concepts tells us a great deal about what would count as an instance of a particular interesting property, such as personhood, knowledge, or rightness. However, an externalist approach to expressions like 'know' and 'right' bypasses the conceptual step in this chain. An externalist about such expressions can say that semantics does a great deal of metaphysical work: to know the semantic properties of a word like 'know' requires knowing its extension (the instances of knowledge) or at least which state is the state of knowing. This externalist does not claim, of course, that such facts are easy to discover and, by definition, she denies that they are available *a priori*, but she can insist that they are semantic and that they are also among the most coveted metaphysical prizes. Thus, such an externalist might urge that one *can* go from semantic premises to substantive conclusions.

The first thing that one might note, in response, is that up to this point the paper has remained neutral regarding one's preferred approach to semantic theory. (It has also attempted to remain neutral, to the extent possible, on one's preferred approach to metaphysics and conceptual analysis.) This externalist argument, however, explicitly depends upon a rather controversial view about semantics. A quick look at linguistics journals reveals that there are a number of active research programmes that adopt quite diverse approaches to semantics. And Chomsky (2000) goes so far as to deny that there is such a thing as referential semantics of the kind that philosophers of language (externalist and internalist alike) typically practise. Even those philosophers who accept that the work of Kripke and Putnam provides important insights are divided on how these should be honoured by semantic theory. In particular descriptivists, who are likely to be friends of conceptual analysis, are as likely to accept those insights as the externalist whose challenge we are considering. So the kind of semantic programme which would bridge the gap between the linguistic and the metaphysical in the way described by our envisaged objector is contentious in a way that leaves it a hostage to

---

define a theoretical role for this folk notion without also thinking that they are the correct semantics of the phrase 'rational requirement' in English. We would be happy with this outlook. However, some philosophers whose primary concern is not language are unclear as to whether they take themselves to be providing semantics for a regimentation of an expression from natural language, in the above sense, or the semantics of an expression in English. A case in point is John Broome's (2007) semantics for 'requirement'.

[33] See Kingsbury and McKeown-Green (2009, 180-181).

methodological fortune. Moreover, despite the current popularity of semantic treatments of normativity, the metaethicists who undertake these projects seem not to be relying on this kind of externalism. Thus, even if the externalist were to win the day in the semantics wars, that would not straightforwardly support the current semantic approaches to metaethics.

More importantly, if the externalist objector is right about semantics, most if not all of the research involved in doing semantic theory just will be research in metaphysics (rather than in semantics). If 'Geoffrey' refers directly to Geoffrey, and if its reference is determined by an appropriate causal link or acquaintance relation with Geoffrey, then there is nothing more to be said that a consideration of linguistic properties can illuminate. All that could possibly be left to consider is the nature of Geoffrey himself. Similarly, and more to the point, if 'practical wisdom' refers to the property of practical wisdom, and if its connection to this property is secured by some appropriate causal story, then further reflection on this issue can proceed only by considering the nature of practical wisdom itself.[34] To be sure, such an investigation might qualify as semantics. But that is hardly interesting. If our concern as metaethicists is with the nature of practical wisdom, goodness, rightness and the like, an externalist can certainly regard this concern as something that is captured by semantics, rightly construed. But the semantics ends up not doing any work. All of the work is done by whatever is left when all of the words have gone away.

7. Conclusion

Moral theorists contribute to linguistic theory as part of their investigation of the normative and semanticists analyse concepts in order to enrich and contextualise their understanding of linguistic meaning, reference and competence – and they should. However, we argued in section 1 that claims about conceptual matters are often represented as, or inferred from, claims about the meaning, reference, or mastery of words, without sufficient justification for this change of emphasis. Sometimes, as with Francois Schroeter on sentimentalism, the result is that a debate is misdescribed and problematic conclusions are then misattributed to other people. At other times, as with Finlay on the meaning of 'ought' and Wedgwood on the relative priority of the good and the right, we find gaps in the key argument, because a semantic result is automatically assumed to yield a conceptual one. In all these cases a basic distinction between issues to do with language and issues to do with our conceptual commitments is overlooked. As we argued in section 2, as philosophers we should be attentive to this kind of distinction. In section 3, we showed an analyst who opts to do semantics cannot ignore the distinction, because there is more to conceptual structure than semantic structure. In section 4 we suggested that inadequate attention is paid by linguistically inclined metaethicists to the fact that data from language use is of primarily pragmatic, rather than semantic, import. Perhaps it seems obvious to some that certain words are the standard vehicles for certain concepts; however, we argued in section 5 that this is not at all obvious. Finally, in section 6 we responded to the challenge that semantics 'properly' understood is externalist and therefore *just is* metaphysics.

When moral theorists employ linguistic arguments, there should be justifications for the role and scope that they assign to those arguments. We have failed to find adequate justifications for much that is linguistic in metaethics and we have identified a number of significant obstacles that would need to be overcome by anyone who wished to treat conceptual analysis as though it were nothing but linguistic analysis.

These findings are equally relevant to other subdisciplines. Many philosophers seeking results that lie beyond the philosophy of language develop what they take to be accounts of the meaning, reference determination, or mastery of linguistic expressions because they hold that

---

[34] See Zagzebski (2004, 42-47).

such theorising contributes directly, centrally, and decisively to their projects. Some epistemologists defend contextualist (DeRose 2005, 2006) or contrastivist (Schaffer 2004) theories primarily by appealing to semantic analyses of the verb 'to know'. To investigate the colours, theorists rely on the semantics of words for colours (Johnston 1992). To discover the nature of meaning, they consider what 'meaning' means (Grice 1957). To discover what actions are, they analyse the logical form of action sentences (Davidson 1968). To investigate natural kinds, they explore the content and reference of kind terms (Barnett 2000). To probe the nature of dispositions, they investigate the truth conditions of subjunctive conditionals (Manley and Wasserman 2008). In doubting whether the idea of following the addition rule is intelligible, they promote scepticism about the meaningfulness of 'plus' (Kripke 1982). To decide whether there are electrons, they ask whether the sentence, 'There are electrons,' is truth-apt and whether it is true (Dummett 1978).

Thus, our discomfort with the entrenched reliance on semantic theorising about ordinary language in metaethics extends to contemporary philosophy quite generally. This is evident even from our brief illustrative discussions of personhood, knowledge, and belief. However, detailed discussion of the problem as it arises in these other areas is a project for another occasion.

References

Anderson, Elizabeth 1993. *Value in Ethics and Economics*, Cambridge, MA: Harvard University Press.

Aune, Bruce 1967. "Hypotheticals and 'Can', Another Look," *Analysis* 27: 191-5. Reprinted in Watson, 1982, 36-41.

Bach, Kent 2002. "Semantic, Pragmatic" in J. Keim Campbell, M. O'Rourke, and D. Shier, eds, *Meaning and Truth,* New York: Seven Bridges Press, 284-292.

Barnett, David 2000. "Is Water Necessarily Identical to H2O?" *Philosophical Studies* 98: 99-112.

Braddon-Mitchell, David and Frank Jackson 2006. *Philosophy of Mind and Cognition*, 2nd ed, Oxford: Wiley-Blackwell.

Broome, John 2007. "Wide or Narrow Scope?" *Mind* 116/462: 359-370.

Carnap, Rudolph 1956. "Empiricism, Semantics, and Ontology," *Meaning and Necessity: A Study in Semantics and Modal Logic*, 2nd ed, Chicago: University of Chicago Press.

Chomsky, Noam 1981. *Lectures on Government and Binding*, Dordrecht: Foris Publications.

Chrisman, Matthew 2011. "'Ought' and Control," *Australasian Journal of Philosophy*, DOI:10.1080/00048402.2011.611151.

Chrisman, Matthew (forthcoming). "On the Meaning of 'Ought'," *Oxford Studies in Metaethics* 7.

Davidson, Donald 1967. "The Logical Form of Action Sentences," in *Essays on Actions and Events*, Oxford: Clarendon Press, 2001, 105-148.

DeRose, Keith 2005. "The Ordinary Language Basis for Contextualism, and the New Invariantism," *Philosophical Quarterly* 55: 172-198.

DeRose, Keith 2006. "'Bamboozled by Our Own Words': Semantic Blindness and some Arguments Against Contextualism," *Philosophy and Phenomenological Research* 73: 316-338.

Donnellan, Keith 1970. "Proper Names and Identifying Descriptions," *Synthese* 21: 335-358.

Dummett, Michael 1978. *Truth and Other Enigmas*, London: Duckworth.

Finlay, Stephen 2009. "Oughts and Ends," *Philosophical Studies* 143: 315-340.

Gert, Joshua 2011. "Naturalistic Metaethics at Half Price," in *New Waves in Metaethics*, Michael Brady, ed. Houndmills, Basingstoke: Palgrave Macmillan, 36-61.

Grice, Paul 1957. "Meaning," *Philosophical Review* 66: 377-388. Reprinted in Grice 1989, 213-223.

Grice, Paul 1975. "Logic and Conversation," in Grice 1989, 22-40.

Grice, Paul 1989. *Studies in the Way of Words*, Cambridge, MA: Harvard University Press.

Higginbotham, James 1985. "On Semantics," *Linguistic Inquiry* 16: 547-593.

Horgan, Terry and Mark Timmons 2006. "Cognitivist Expressivism," in *Metaethics after Moore*, Terry Horgan and Mark Timmons, eds, Oxford: Oxford University Press, 255-298.

Jackson, Frank 1998. *From Metaphysics to Ethics*, Oxford: Oxford University Press.

Jackson, Frank and Phillip Pettit 1995. "Moral Functionalism and Moral Motivation," *Philosophical Quarterly* 45: 20-40.

Johnston, Mark 1992. "How to Speak of the Colors," *Philosophical Studies* 68: 221-263.

Kaplan, David 1989. "Demonstratives: An Essay on the Semantics, Logic, Metaphysics, and Epistemology of Demonstratives and Other Indexicals," in *Themes from Kaplan*, J. Almog, J. Perry, and H. Wettstein, eds, New York: Oxford University Press.

Kingsbury, Justine and Jonathan McKeown-Green 2009. "Jackson's Armchair: The Only Chair in Town?" in *Conceptual Analysis and Philosophical Naturalism*, David Braddon-Mitchell and Robert Nola, eds, Cambridge, MA: MIT Press, 159-182.

Kratzer, Angelika 1981. "The Notional Category of Modality," in *Words, Worlds, and Contexts*, Eikmeyer and Reiser, eds, Berlin: de Gruyter, 38-74.

Kripke, Saul 1972. *Naming and Necessity*, Cambridge, MA: Harvard University Press.

Kripke, Saul 1982. *Wittgenstein on Rules and Private Language*, Cambridge, MA: Harvard University Press.

Lehrer, Keith 1968. "Cans without Ifs," *Analysis* 29: 29-32. Reprinted in Watson 1982, 41-45.

Ludlow, Peter 2003. "Referential Semantics for I-Languages," in *Chomsky and His Critics*, Louise Antony and Norbert Hornstein, eds, Malden, MA: Blackwell Publishing, 140-161.

Lynch, Michael 2005. "Alethic Functionalism and Our Folk Theory of Truth," *Synthese* 145: 29-43.

Manley, David and Ryan Wasserman 2008. "On Linking Dispositions and Conditionals," *Mind* 117/465: 58-84.

Margolis, Eric and Stephen Laurence, "Concepts", *The Stanford Encyclopedia of Philosophy (Fall 2011 Edition)*, Edward N. Zalta, ed., http://plato.stanford.edu/archives/fall2011/entries/concepts/.

McDowell, John 1987. "Projection and Truth in Ethics," reprinted in *Mind, Value, and Reality*. Cambridge, MA: Harvard University Press, 151-166.

Putnam, Hilary 1978. "The Meaning of 'Meaning'," in *Mind, Language, and Reality*, Cambridge: Cambridge University Press, 215-271.

Raz, Joseph 2001. "Two Views of the Nature of the Theory of Law: A Partial Comparison," in *Hart's Postscript: Essays on the Postscript to The Concept of Law*, Jules Coleman, ed, Oxford: Oxford University Press, 1-37.

Schaffer, Jonathan 2004. "From Contextualism to Contrastivism," *Philosophical Studies* 119: 73-103.

Schroeder, Mark 2004. "The Scope of Instrumental Reason," *Philosophical Perspectives* 18: 337-364.

Schroeter, François 2006. "The Limits of Sentimentalism," *Ethics* 116: 337-361.

Schroeter, François and Laura Schroeter 2009. "A Third Way in Metaethics," *Nous* 43: 1-30.

Schroeter, Laura 2004. "The Limits of Conceptual Analysis," *Pacific Philosophical Quarterly* 85: 425-453.

Smith, Michael 1994. *The Moral Problem*, Oxford: Blackwell Publishing.

Sturgeon, Scott 2007. "Normative Judgement," *Philosophical Perspectives* 21: 569-587.

Watson, Gary, ed 1982. *Free Will*, Oxford: Oxford University Press.

Wedgwood, Ralph 2001. "Conceptual Role Semantics for Moral Terms," *Philosophical Review* 110: 1-30.

Wedgwood, Ralph 2006. "The Meaning of 'Ought'," *Oxford Studies in Metaethics* 1: 127-160.

Wedgwood, Ralph 2009. "The 'Good' and the 'Right' Revisited," *Philosophical Perspectives* 23: 499-519.

Wiggins, David 1987. *Needs, Values, Truth*, 3rd ed. Oxford: Oxford University Press.

Zagzebski, Linda 2004. *Divine Motivation Theory*, Cambridge: Cambridge University Press.