Stephen Finlay, *Confusion of Tongues: A Theory of Normative Language*.
Oxford: Oxford University Press, 2014. Pp. 278. $65.00 (cloth).

Reviewed by Daniel Fogal, Uppsala University

Stephen Finlay's *Confusion of Tongues* is a bold and sophisticated book. The overarching goal is metaphysical: to reductively analyze normative facts, properties, and relations in terms of non-normative facts, properties, and relations. But the method is linguistic: to first provide a reductive analysis of the corresponding bits of normative language, with a particular focus on 'good', 'ought', and 'reason'. The gap between language and reality is then bridged by taking linguistic analysis as a guide to conceptual analysis, and conceptual analysis as a guide to metaphysical analysis. Finlay thus assumes that words are semantically associated with concepts, and that the composition of concepts reveals the "metaphysical composition or essence" of the corresponding properties (6).

Finlay is aware that his methodology might seem at best naïve and at worst gravely mistaken. He addresses some of the most obvious concerns up front, but his main defense of the "analytic method" consists in its execution over the course of the book. Even if it is doubtful he succeeds—some seeds of doubt will be sowed below—there is nonetheless much to be gained from Finlay's dense and detailed discussion. I therefore highly recommend it to any philosopher interested in the semantics and pragmatics of normative language.

The book divides into two parts, with Chapter 1 setting the stage. The first part defends an "end-relational" account of 'good' (Chapter 2), 'ought' (Chapter 3), and 'reason' (Chapter 4). Very roughly, for something to be good is for it to *increase the probability* of (and hence be "good for") some contextually relevant end or outcome, for it to be that one ought to do something—call it φ—is for φ-ing to be *most probable* (and hence "best") given that the relevant end obtains, and for something to be a reason to φ is for it to be an *explanation why* it would be good (in some way, to some degree) to φ. These analyses, if successful, motivate the idea that *all* normative locutions can be analysed either directly or indirectly in terms of "probabilistic relations in which things stand to particular 'ends' or potential states of affairs that vary from context to context" (1). Normative claims would thus boil down to claims about what probabilifies what (and why), and normative facts would reduce to probabilistic facts (and what explains them).

Although Finlay's end-relational semantics is consistent with many ordinary uses of 'good', 'ought', and 'reason', a pressing worry is that it fails to accommodate the full range of such uses, including distinctively normative ones. This sets the agenda for the second half of the book, in which Finlay makes heavy use of pragmatic principles in an effort to explain "the distinctive features of peculiarly moral and deliberative uses" of normative language (1). The challenges include the

close connection such uses have with motivation (Chapter 5), the existence and weighing of multiple ends (Chapter 6), the apparently categorical and non-relational nature of moral claims and claims about intrinsic value (Chapter 7), and the possibility of substantive disagreement (Chapter 8). Finlay then concludes by assessing the overall merits of the resulting metanormative view (Chapter 9), arguing that it provides "a maximally simple and conservative explanation" across all relevant dimensions of evaluation: "linguistically, metaphysically, psychologically, epistemologically, and ethically" (246).

A major theme throughout the book—and the inspiration for the title—is that most normative and metanormative disputes result from a "confusion of tongues". More specifically, Finlay argues that most fundamental *normative* disagreements—over what is good, what one ought to do, and so forth—are cases in which we are "talking past each other, using the same sentences and words to refer to different facts and properties", and most *metanormative* disagreements—over the nature of normative facts, properties, judgments, and so forth—arise because of "a failure to understand our own use of normative language" (1).

The audacity of Finlay's project is part of what makes it so rich and interesting, and yet also so predictably hard to defend. In what follows I'll consider three challenges to Finlay's project that deserve more attention than they receive in the book. The first concerns Finlay's claim to have provided a *reductive* theory of normative language, the second concerns his claim to have provided a *unified* theory of normative language, and the third concerns his claim to have provided a *correct* theory of normative language.

The first challenge isn't really a challenge, but instead a complication concerning Finlay's claim to have provided distinctively linguistic or semantic support for a form of reductive naturalism. For although Finlay's end-relational *metanormative theory* is indeed reductive, with normative facts reducing to probabilistic (and hence "natural") facts, his end-relational *semantics* is not. And that's because the semantics itself is silent concerning the status of the ends implicated in our use of 'good', 'ought', and 'reason(s)'. So even if normative language is in fact end-relational, it may be that some ends are themselves normative. In principle, then, a non-naturalist could accept Finlay's semantic analyses while also insisting on the existence of certain irreducibly normative ends or outcomes—for example, that justice is served, that a person's autonomy is respected, or that a virtuous character trait is developed. Indeed, this would provide a straightforward way of distinguishing distinctively *normative* uses of 'good', 'ought', and 'reason' from the rest—distinctively normative uses would be those that involve irreducibly normative ends. Why does a parent's use of 'good' in 'It's good that you told the truth' strike us as normatively robust in a way that a thief's use of 'good' in 'It's good that we didn't get caught' does not? The non-naturalist who helps themselves to Finlay's semantics has a ready answer: because the end that's relevant to the former (e.g. doing what's morally right) is normative and the end that's relevant to the latter (e.g. succeeding in one's mischief) is not.

Finlay's argument in favor of reductive naturalism thus rests on a crucial metaphysical assumption: that all ends implicated in our use of normative language

are ultimately non-normative. Although Finlay recognises this point, he doesn't emphasize it as much as one might expect, especially since it complicates his avowed goal of providing distinctively *linguistic* or *semantic* support for reductive naturalism. Indeed, Finlay implores the reader "not to forget that the simplicity, conservatism, and predictive power *of this semantics* constitutes my primary evidence that the end-relational theory is not only consistent with our practices of normative speech and thought, but also correct" (18, my emphasis; cf. p. 246). But this claim and others like it are confusing, since it's only in combination with the *extra*-semantic assumption that all ends are non-normative that Finlay's end-relational semantics supports his end-relational—and reductive—metanormative theory. Finlay himself, however, uses 'the end-relational theory' and 'the end-relational semantics' interchangeably, and writes as if the extra-semantic assumption is part of the semantics itself. This terminological unclarity is unfortunate, since it wrongly encourages the thought that a form of naturalism falls out of Finlay's semantics rather than being an additional ingredient from the start. As a result, even a careful reader could easily be forgiven for thinking that Finlay is trying to pull a metaphysical rabbit out of an end-relational semantic hat (as Laskowski 2014 aptly puts the point).

It doesn't follow that Finlay's semantics is metanormatively irrelevant, however. On the contrary, it plays a central role in motivating Finlay's particular metanormative view *vis-á-vis* its naturalistic competitors, the likes of which include expressivism, prescriptivism, and "synthetic" reductive naturalism. So Finlay's semantics *together with the assumption of naturalism* is what motivates his *particular brand* of naturalism—i.e. analytic reductive naturalism. And that's no mean feat. But it doesn't motivate naturalism itself. That's a separate issue, and one to be settled on more traditional, non-linguistic grounds.

The second challenge concerns the simplicity and uniformity of Finlay's semantics, rather than its metaphysical underpinnings. In outlining the case for his view, Finlay writes:

> In particular, I'll pursue the desideratum of a maximally simple *semantics*. Whereas metaethicists have generally been bad linguists, offering definitions of normative words based on narrow focus on particular kinds of use, these words are used in a much broader range of ways, which commits rival theories to extensive lexical ambiguities. But by following the clues from language itself we'll identify single, shared, unifying meanings for these words that accommodate virtually all their uses. (16)

Although Finlay rightly emphasizes the wide variety of ways in which words like 'good', 'ought', and 'reason' are used, and rightly insists that positing extensive lexical ambiguities is generally a cost, it's important to recognize that not all ambiguities are created equal. It's standard, for example, to distinguish two forms of ambiguity: polysemy and homonymy. Roughly put, a word form is polysemous just in case it is associated with two or more distinct but related meanings, while a word form is homonymous just in case it is associated with two or more unrelated

meanings. Ambiguity in both its forms (polysemy and homonymy) is thus to be contrasted with univocality (or monosemy), where a word is associated with only one meaning.

A textbook example of homonymy is 'bank'—compare 'I sat on the bank of the river' with 'I deposited money at the bank'—while standard cases of polysemy include 'line' and 'see':

> I drew a *line*; She read a *line*; He has *lines* around his eyes; Clothes hung on a *line*; Jorge waited in a *line*; I made a *line* of bad decisions (cf. Falkum and Vincente 2015).

> Did you *see* the sunset?; I *see* your point; *See* how it sounds; You should *see* a doctor; *See* that you don't break it; Sam's been *seeing* Maxine (cf. Pullum 2016).

Suffice it to say, polysemy is utterly pervasive in natural language—much more so than homonymy—and it affects both content and function words. It also tends to be both systematic and productive, with similar patterns of polysemy applying to similar words across many languages.

Although the existence and prevalence of polysemy raises a host of difficult theoretical and empirical questions—see Falkum and Vincente (op. cit.) for an overview—the fact remains that if we follow "the clues from language itself" then we should *not* expect many, if any, ordinary words to have a "single, shared, unifying" meaning. Of course some words are more polysemous than others, but with well-worn and widely used words such as 'good' and 'reason' the prospects of a unified analysis are particularly bleak. Indeed, given that so many words of far more recent vintage and far less widespread use are polysemous, and significantly so, it would border on the miraculous if 'good' or 'reason' managed to avoid the same fate.

To be fair, Finlay is aware of the possibility that 'good' and 'reason' are polysemous, and concedes that there may be "some outlying polysemy or idioms" that his semantic analyses don't capture (35; see also p. 86). Nonetheless, Finlay doesn't seem to appreciate just how pervasive polysemy is, and how much of a difference this makes to the overall dialectic. The assumption that we should favor highly unified analyses of words like 'good', 'ought' and 'reason' is important for Finlay's purposes, since it serves to both motivate his own, highly unified—and at times seemingly ad hoc—analyses as well as discount the possibility of there being other, less unified analyses that fit the data more naturally. Finlay recognises the presumption of semantic simplicity and uniformity is defeasible, but he nonetheless overestimates its plausibility, since for pretty much all ordinary words the presumption isn't merely defeasible, it's positively defeated. It's therefore a mistake to think that the "default hypothesis should be that 'good' [or 'reason'] has a single, unified semantics" (19). Instead, the default hypothesis for pretty much any ordinary noun, adjective, or verb should be that it is polysemous, and thoroughly so.

To take just one example relevant to Finlay's purposes: it's standard to distinguish between the reasons *why* something is the case (commonly called "explanatory reasons"), the reasons *for which* or *on the basis of which* someone does something ("motivating reasons"), and the reasons for someone *to* do something ("normative reasons"). There are a range of locutions associated with each. Here are three:

> **Explanatory (ER):**   *r* is a (or the) reason (why) *q*.
> **Motivating (MR):**   *S*'s reason for φ-ing is that *q*.
> **Normative (NR):**   *r* is a reason (for *S*) to φ.

As Finlay notes, in drawing such distinctions it is often assumed—if only implicitly—that (ER), (MR), and (NR) each involve different (though presumably related) senses of 'reason'.

Finlay himself demurs, however, proposing that in each case 'reason' simply means *explanation why*. Although tailor-made for (ER), this proposal doesn't obviously fit (MR) and (NR)—substituting 'explanation why' for 'reason' in either results in a sentence form of dubious coherence (as well as grammaticality). Finlay nonetheless goes to great lengths to secure uniformity, making generous appeal to ellipsis. In the (near) final analysis of (NR), for example, his account predicts that a "fully explicit (pro tanto) normative reason sentence can take the [following] grammatical form":

> **(NR*)**   *r* is a reason/explanation relative to [background information] b_1 why it is good for [some end] *e*, given b_2, for *s* to φ. (97)

Finlay also offers a couple of similarly complex analyses of (MR) sentences, without deciding between them (106-114). Importantly, however, neither analysis accounts for closely related sentences in which the 'that'-clause ('that p') is replaced by an infinitive ('to ψ'):

> **(MR-*to*):**   *S*'s reason for φ-ing is to ψ.

Sentences of this form pose a challenge to Finlay's unified analysis of 'reason'. That's because (MR-to) sentences don't seem to be factive—Rachel's reason for studying law can be to make money, for example, even if she doesn't end up making any money. According to Finlay, however, (MR) sentences *are* factive—if Rachel's reason for going into law is that it pays well, for instance, it must be the case that going into law does in fact pay well. Sentences of the form (ER) and (NR)—if Finlay's right—are also factive: the fact that Jim quit his job can't be a reason he's broke if it turns out he's not actually broke, and the fact that it's raining can't be a reason (why it's good) to bring an umbrella unless it is in fact good to bring an umbrella.

The factivity of (ER), (MR), and (NR) is important to Finlay because he claims that reasons are explanations why, and that 'explanation why' is "doubly factive: *r*

can only be an explanation why *p* if both *r* and *p* are true" (89). But as we've seen, 'reason' in (MR-*to*) doesn't seem to be factive in this way. This is a problem for Finlay, especially given the obvious similarity between the use of 'reason' in (MR-*to*) and the use of 'reason' in (MR)—if the former eludes Finlay's reach we should expect the latter to as well.

This is just one instance of the apparent polysemy of 'reason'. I'll provide another below. The general point, however, is that with respect to both (NR) and (MR) sentences Finlay purchases uniformity at the level of lexical semantics ('reason') at the cost of considerable—and to a large extent grammatically unmotivated—complexity at the level of logical form (e.g. (NR*)). Although elements of his analyses are independently motivated, the driving force remains the presumption of semantic uniformity. Without it Finlay's unified analysis of 'reason' loses whatever lustre it might have otherwise enjoyed, and competing theories that posit polysemy look considerably more promising.

In responding to the challenge from polysemy, it wouldn't be totally unreasonable for Finlay to stick to his guns and continue to insist on the correctness of his analyses (perhaps modulo (MR)/(MR-*to*)). For even if we might not have *expected* there to be highly unified analyses of 'good', 'ought', and 'reason', and even if the assumption of semantic uniformity is empirically unwarranted, perhaps—thanks to Finlay—we're fortunate enough to be in a position to provide them. This brings us to the third challenge, which is that there is a distinctively *normative* use of 'reason' that Finlay fails to capture. (The same is true of 'good' as it appears in sentences like 'Pleasure is good'—for details, see Skarsaune 2015.) This problem is particularly pressing, since it poses a direct threat to Finlay's form of analytic reductive naturalism.

It's worth noting that 'reason' in its normative sense exhibits what Apresjan (1974) calls *regular* polysemy. According to Apresjan, the polysemy of a word A with the meanings a_1 and a_2 is regular (or systematic) "if, in the given language, there exists at least one other word B with the meanings [b_1 and b_2], which are semantically distinguished from each other in exactly the same way" as a_1 and a_2 (op. cit., 16). The polysemy of A is *irregular* (or non-systematic) if the difference between a_1 and a_2 "is not exemplified in any other word of the given language" (op. cit., 16). Clear cases of regular polysemy include nouns that alternatively denote containers and amounts contained (e.g. 'cup', 'bottle') as well as nouns that alternatively denote institutions and the places they reside (e.g. 'school', 'church'), but there are many others. (For an overview, see Dölling forthcoming.)

The particular pattern of regular polysemy I wish to focus on concerns the count/mass distinction, since 'reason' in its normative—but not explanatory or motivating—sense is standardly used both as a count noun ('Julie has many *reasons* to lie') and as a mass noun ('Julie has lots of *reason* to lie'), and Finlay falters in his analysis of the latter. Intuitively, count nouns denote (classes of) "things" that are countable, and hence can occur with cardinal numerals ('one', 'two', 'three'...) and take plural form (-*s*), while mass nouns denote "stuff" that's not countable, and hence do not occur with cardinal numerals and are generally singular or unmarked. There's a lot to be said about the mass/count distinction, but what matters most is

that 'reason' is regularly used both ways. The relationship between the two uses has nonetheless been neglected, with the vast majority of normative theorists focusing exclusively on 'reason' as a count noun, despite also using it as a mass noun. (Finlay is an admirable exception; see below.) Claims concerning what one has *most* or *more* reason to do, for example, are ubiquitous and obviously not equivalent to the corresponding superlative and comparative claims concerning what one has most or more *reasons* to do.

This neglect is surprising since it is extremely common for nouns in languages with mass-count syntax (such as English) to be used both ways, and there are a number of well-established patterns underlying such alternations. This raises the obvious question: what about the relationship between mass and count uses of 'reason'? Although space precludes detailed discussion, I argue elsewhere (Fogal, 2016) that the relationship is the same as that which holds between mass and count uses of 'pleasure', 'sorrow', and 'light'. Just as pleasures are sources of pleasure, sorrows are sources of sorrow, and lights are sources of light, so (normative) reasons are "sources" of reason. Importantly, however, there's an explanatory asymmetry involved: in each case the things denoted by the count noun (pleasures, lights, reasons) are understood in terms of the role they play in generating or explaining the "stuff" denoted by the mass noun (pleasure, light, reason), rather than vice versa. So just as it would be a mistake to analyze pleasure in terms of pleasures, or light in terms of lights, so it would be a mistake to analyze reason in terms of reasons. Instead, we should analyze (normative) reasons in terms of reason: reasons to φ are things which (help) explain why there is reason to φ. This proposal straightforwardly predicts the intuitive equivalence of (R) and (R')-(R''):

> (R)   The dark clouds approaching <u>are</u> <u>a reason</u> to think it will rain.
> (R')   The dark clouds approaching <u>give</u> us <u>reason</u> to think it will rain.
> (R'')   <u>There is reason</u> to believe it will rain <u>because</u> there are dark clouds approaching.

All of this poses a threat to Finlay's unified analysis of 'reason' as meaning *explanation why*. For although Finlay recognizes that 'reason' is used as a mass noun, he downplays its significance and attempts to unify it with his analysis of the count noun. In particular, he offers the following analyses of the comparative ('more reason') and superlative ('most reason') constructions:

> **More Reason**:   'There's more reason to φ than to ψ' means *there's an explanation why it's more good (better) to φ than to ψ*.
> **Most Reason**:   'There's most reason to φ' means *there's an explanation why it's most good (best) to φ*. (91)

Finlay realizes that these analyses "aren't fully compositional", but thinks this is excusable because "the expressions are evidently idiomatic" (92). The latter claim is plainly false. As the examples above make clear, the relationship between the

mass noun and count noun is systematic (in Apresjan's sense), and sentences involving 'more/most reason' are no more idiomatic than those involving 'more/most pleasure' and 'more/most light'.

We thus have ample reason to think 'reason' is polysemous, not just as it appears in (MR)/(MR-*to*) but also as it appears in (NR). Although Finlay is right that the various uses of 'reason' as a count noun are all explanatory—to be a reason (of any kind) is, in part, to be something that helps explain something else—he's wrong to think there are no further differences between them and he's wrong to think the use of 'reason' *as a mass noun* is also explanatory. On the contrary, the mass noun seems to be inherently normative, which would explain why (in contrast to the count noun) there are no merely "explanatory" or "motivating" uses of it.

Although I don't take any of the challenges outlined above to be decisive, taken together they do cast doubt on the ability of Finlay's end-relational semantics to deliver the desired metanormative goods. No matter the ultimate fate of his view, however, Finlay is to be applauded for his ambition and for consistently arguing with such clarity and care. Finlay's admirable attention to detail and keen eye for subtleties make *Confusion of Tongues* an invaluable contribution to the existing literature, and one that will continue to shape and inform a number of important debates going forward.[1]

References

Juri D. Apresjan, "Regular Polysemy," *Linguistics* 12, no. 142 [1974]: 5–32.

Johannes Dölling, "Systematic Polysemy," in *The Blackwell Companion to Semantics*, eds. Lisa Matthewson, Cécile Meier, Hotze Rullman, and Thomas Ede Zimmermann [Hoboken, NJ: Blackwell, forthcoming].

Ingrid Falkum and Agustin Vincente, "Polysemy: Current Perspectives and Approaches," *Lingua* 157 [2015]: 1–16.

Daniel Fogal, "Reasons, Reason, and Context," in *Weighing Reasons* [Oxford: Oxford University Press, 2016]: 74-103.

Nicholas Laskowski, "How to Pull a Metaphysical Rabbit out of an End-Relational Semantic Hat," *Res Philosophica* 91, no. 4 [2014]: 589–607.

Geoffrey Pullum, "Polysemy and Maturity," *Lingua Franca* [February 19, 2016]: http://chronicle.com/blogs/linguafranca/2016/02/19/polysemy-and-maturity/

Knut Olav Skarsaune, "How To Be a Moral Platonist," *Oxford Studies in Metaethics* 10 [2015]: 245-272.