Kenesei
Research Institute for Linguistics
Hungarian Academy of Sciences Pf. 360
1394 Budapest, Hungary
[kenesei.istvan@nytud.mta.hu]

Lipták
Leiden University Centre for Linguistics
Leiden University
P.O. Box 9515
2300 RA Leiden, The Netherlands
[A.Liptak@hum.leidenuniv.nl]

Confusion of tongues: A theory of normative language. By Stephen Finlay. Oxford: Oxford University Press, 2014. Pp. 288. ISBN 9780199347490. $69 (Hb).

Reviewed by James Lenman, *University of Sheffield*

We begin with *good*. *Good*, Finlay suggests, is an incomplete predicate. When we say something is good we are saying something of this form.

(1) It is good for e if p.

Where p is some proposition (for *good*, he argues, is fundamentally expressive of a propositional operator), and e is some end salient in the context. Such salience is a slippery thing. Sometimes *Drinking a lot of gin is good* might be relative to the end of having fun. But we can still ask *Drinking a lot of gin may be fun, but is it good?* to raise the question of its conduciveness to some other end comparable in eligibility for salience, such as health perhaps. So we have a form of reductive naturalism that offers to defuse the OPEN QUESTION ARGUMENT. Or at least we do when we complete the reduction in the way F proposes.

(2) 'It is good for e if p' means that p increases the probability of e.

Or more precisely:

(3) 'It is good for e if p' means that pr(e/p & b) > pr(e/~p & b).

Where b specifies certain background conditions.

Along the way in discussing *good*, F has a swipe at the Aristotelian moral functionalist. A sentence saying *X is a good K*, where K is some functional kind, has both a functional reading, where the end served by the state of affairs denoted by p in which X features as a K is that end e_K specified by the function of K; and a nonfunctional reading, where X is something that happens to be a K featuring in some state of affairs good for some contextually salient e. F now claims we can distinguish these readings in the following way. On a nonfunctional reading, X being a good (or bad) K entails X being a K. On a functional reading, it does not. (This test is credited to Shyam Nair.) To be a good (or bad) person, he then proposes, entails being a person. Likewise, to be a good (or bad) human being entails being a good human being. Goodbye moral functionalism. It would be good to see this very quick argument slowed down a bit. As it stands, I am not quite clear what the rules would be for seeking counterexamples. An ice-axe is not a weapon, but it may be a good weapon. Only, if it is a good weapon, it surely is a weapon. Do we take *weapon* here to mean something made for a martial purpose or something serviceable to one? On the latter understanding, *weapon* surely still counts as a functional kind, but it is hard here to see how a good weapon could fail to be a weapon. Maybe things get clearer if we accentuate the negative. A postage stamp is a really bad weapon. Indeed, it is (partly for that reason) not a weapon at all.

We turn next to *ought*. F focuses on the instrumental conditional.

(4) If Max wants to evade arrest, then he ought to mingle with the crowd.

F tackles this in two steps. The first offers an analysis of the more straightforward:

(5) If Max is going to evade arrest, then he has to mingle with the crowd.

This F takes to mean that, in all possibilities consistent with certain background conditions in which Max will evade arrest, Max will mingle with the crowd. F then offers the neat proposal that we read the original *wants to* as a relevance or 'biscuit' conditional akin to 6, which is supposedly elliptical for 7.

> (6) If you want biscuits, there are some on the table.
>
> (7) If you want biscuits, [then this will be relevant information:] there are some on the table.

Likewise, for 8 we get the preliminary analysis proposed in 9.

> (8) If Max wants to evade arrest, then he ought to mingle with the crowd.
>
> (9) If Max wants to evade arrest, [then this will be relevant information: If Max is going to evade arrest,] then he has to mingle with the crowd.

This is preliminary because while something like the foregoing is okay with the sort of strong necessity expressed by *must*, it is less clearly what we want in order to capture the weaker necessity of *ought*. With *ought*, F suggests, the relevant information is that mingling with the crowd is Max's likeliest way (with respect to some background of alternatives) to evade arrest. Having handled hypothetical *ought*-claims in this way, F suggests we handle them all this way, taking nonhypothetical *ought*s as elliptical for hypothetical *ought*s directed at contextually salient ends.

The last of the three key normative concepts dealt with here is 'reason'. Reasons, says F, are explanations of what makes it the case that something is good. A reason for me to go the airport is an explanation of why it is good if I go to the airport.

This 'end-relational' view is a wholly descriptivist, modestly naturalistic account of the semantics of normative language. It might thus invite the charge, often leveled at such accounts, of failing to do justice to the so-called practical significance of such language that is naturally taken both as expressing the desires or intentions of the speaker and as functioning to influence those of the auditor. But in fact, F argues, his end-relation theory accounts for this practicality and does so better than rival expressivist views that might be thought better suited to the task. For F, the practical import of normative utterances (and judgments) is accounted for at the level of pragmatics, not semantics. Where the end e is not explicitly stated, it is made salient by context in various ways, but it is natural on pragmatic grounds to default to the assumption that e is the speaker's 'conversational end', whatever she is concerned to accomplish that motivates her to speak in the first place. And in cooperative contexts, where there is a conversational end shared by speaker and auditor, it is again natural to default to the assumption that e is this end. And where this is the case the speaker's utterance amounts to a recommendation expected to influence the desires and intentions of the auditor.

F's view is thus, as he puts it, 'quasi-expressivist'. It allows that normative utterance typically expresses desires or other desire-like states of the speaker, but, unlike expressivism, it takes this to be an aspect not of the semantics of such utterances but of their pragmatics. We should reject expressivism, F argues, first because: it 'commits us to an uncompromising rejection of the first horn of the Euthyphro dilemma: that we desire or are motivated toward things because we judge them to be good rather than vice versa' (141); second, because it compels us to understand normative thought as a matter of arbitrary choice; and hence, third, because it leaves nothing for moral persuasion to be except a form of nonrational psychological manipulation where the preferences expressed in the respective utterances are in conflict; and finally because it fails to make sense of the recommendatory force of normative utterance: why should my auditor care about the attitudes I express? These are interesting complaints, but I do not know how much real dialectical advantage they secure. The first three might be thought to look no more dangerous to expressivism than to a view like F's own, which also seems ultimately to leave our choice of final ends unsupported other than by the brute contingencies of what we find ourselves stably disposed to desire. The last is an interesting worry. *You ought to get that tire changed*, I say to you, thereby, according to the expressivist, expressing an attitude of mine. But why would that be of any interest to you? Well, it might very well not be. But if it is not it would be, at least ordinarily, a pointless thing for me to say, so plausibly my saying it ordinarily presupposes that your attitudes (here toward e.g. the safety of the car you drive) are sufficiently on the same page as mine for the con-

siderations that here speak to my concerns in providing my utterance with reasons also to speak to yours. Of course, that is an appeal, rather in the spirit of F himself, to the pragmatics of normative utterance, but there is no reason why, if the end-relational theorist can address so many problems with such appeals, the expressivist cannot do the same.

Matters become complicated for F's theory when we have, as we of course very often do, MANY ends that are relevant to the choices we employ normative judgment to make. Often our normative concern is to balance these ends and determine what we should do 'all things considered'. Here the salient end for whoever is the salient person—speaker, auditor, agent—will be, F argues, 'the attainable conjunction of … [her] desired ends that she most prefers' (158). Things get more complicated still when we throw uncertainty into the mix, at which point F suggests that the default end will be whatever maximizes the expected utility of the salient person, though this has to be read de re rather than de dicto, since, of course, no sane person is only interested in whatever maximizes her utility AS SUCH. It is natural to wonder how this account, already to my philosophical ear a little strained at this point, is to extend to those cooperative cases where the salient persons are PLURAL. Do we then take the salient end as their aggregate utilities? But that would seem to build in too much substantive normative theory, and it is hard to see what alternative proposal would avoid doing the same.

But don't we think there are CATEGORICAL normative demands? Certainly we make demands on each other, F suggests, that have what he calls 'pragmatic categoricity'. Here we address each other with normative claims without regard for whether they have any desire for the relevant end. In general, F suggests, normative claims where we fail to explicitly specify the relevant end pragmatically presuppose that we are in a cooperative context where the end is shared. In cases of categorical use, however, this presupposition is false, but we speak as if it were not, so as to endow our utterance with a kind of 'rhetorical objectivity' that puts maximal pressure on the auditor to align her motivation with what we are demanding of them. The last substantial chapter (before a brief concluding one) offers an account of normative disagreement that again leans heavily on pragmatic considerations. F argues that there can be normative disagreement without semantic inconsistency where the desires that his 'quasi-expressivism' allows to be expressed by normative judgments are in conflict.

F's book is a hard and intricate read that will tax any reader. However, it is abundantly worth the effort, easily living up to the very high expectations that those of us who have followed F's earlier writings in the journal literature are bound to bring to it. This is one of the richest, most sophisticated, and most impressive books on metaethics to have been published in my lifetime. Everyone with any interest in normative language ought to read it. Those who would seek to defend reductive naturalist views of the sort F develops here will find it a treasure trove of dialectical resources that they will want to plunder repeatedly. Those who seek to attack such views or to defend rival views will find it a challenge it would be shameful to ignore.

Department of Philosophy
University of Sheffield
45 Victoria St
Sheffield 53 7BQ, United Kingdom
[j.lenman@sheffield.ac.uk]

Automaton theories of human sentence comprehension. By JOHN T. HALE. (CSLI studies in computational linguistics.) Stanford, CA: CSLI Publications, 2014. Pp. 204. ISBN 9781575867472. $27.50.

Reviewed by SASHANK VARMA, *University of Minnesota*

Linguistics and psychology have been intertwined since at least the beginning of the cognitive revolution. In the early 1950s, the psychologist George Miller was busy applying mathematical