

## Reply to Worsnip, Dowell, and Koehn (with *Précis of Confusion of Tongues*)

Stephen Finlay

This is the author's preprint version of an article published in *Analysis* 80(1), 2020: 99-101, 131-147.

### *Précis*

In June 2019, the Canadian Philosophical Association hosted a symposium on my book *Confusion of Tongues: A Theory of Normative Language* (OUP 2014), with critical commentaries from Glen Koehn (the organizer), Janice Dowell, and Alex Worsnip. We are grateful to *Analysis Reviews* for publishing the final version of this exchange of ideas. This précis summarizes COT's central ideas in order to frame the debate for readers.

The guiding idea of *Confusion of Tongues* (COT), and the main point of the title, is that the central puzzles of metaethics—concerning inter alia the nature of normative facts and properties, our epistemic access to them, how they motivate us, and why they have authority over us—stem from our failure to reflectively understand our own ordinary first-order practice of normative discourse. This idea motivates COT's method of conceptual/semantic analysis, and the book as a whole attempts to demonstrate that solutions to these central puzzles can largely be found by attending to the clues provided by natural language. The metaethical theory that emerges from this investigation purports to be a form of descriptivism and analytic reductive naturalism, but with a significant “quasi-expressivist” component. It advocates a form of Humean/Nietzschean perspectivism on normative thought, giving a mixed verdict on the classic Euthyphro question: while the *derivatively* good is desired because it is judged to be good, things are judged *finally* good because they are desired.

My adoption and defense of an analytic method rejects a widely held view that such an approach is futile in metaethics. A secondary theme of COT is that failure to successfully analyze our normative concepts, as well as puzzlement over normative thought and talk, is due largely to failure to reflectively understand the complex pragmatics of how we use this language in context in pursuit of our conversational ends. Accordingly, COT has two parts. In the first part, I investigate the semantics of the central normative terms ‘good’ (Ch. 2), ‘ought’ (Ch. 3), and ‘reason’ (Ch. 4), by seeking a unifying theory that can accommodate as far as possible all the different ways we use them—focusing on their nonmoral and even non-normative uses but temporarily setting aside the moral and practical uses of special interest to philosophers which generate metaethical puzzles. In the second part, I argue that this semantics provides solutions to those puzzles once combined with a single, intuitive basic principle of pragmatics.

What emerges from the earlier chapters on semantics is a form of relativist and contextualist analysis I call an *end-relational theory*: semantically, normative language is descriptive language about relations in which things stand to various *ends*. Ends are potential states of affairs or outcomes, the kind of thing that can be contents of desires, hopes, and intentions, but they are not those psychological attitudes themselves. Sometimes these ends are made explicit, as in talk about what one ought to do *in order that e*, or about what is good *for e*, but at other times they are left implicit and assumed, and have to be

identified as salient in the context. To be “good” for an end *e*, I argue, is just to raise the probability of *e*, and what “ought” to be done in order that *e* is just whatever most raises the probability of *e*. I analyze talk about “reasons” generally as being about explanations, and talk about specifically normative reasons to do A as being about explanations why it would be good (relative to some salient end *e*) to do A.

The end-relational theory is a member of a familiar family of instrumentalist theories about normative properties. Philosophers have largely discounted such theories on the grounds that they are unable to account for central features of specifically moral and practical thought and talk. COT’s later chapters on pragmatics attempt to refute these objections, by demonstrating that the end-relational theory actually predicts all these features once we take into consideration the pragmatics of using end-relational language in the relevant contexts. In Chapter 5 I first argue for a fundamental principle of pragmatics from which all my subsequent solutions are then derived: roughly, that we always speak in the way that we believe best for our conversational ends. I then address a first metaethical problem concerning “practicality”, or the close connection between bare/unrelativized normative utterances and corresponding motivation (“motivational internalism”). I argue that if the end-relational theory is correct, then uses of normative language without explicit relativization must be implicitly relativized to a salient end, and that in general, ends are salient when they are the desired/preferred ends of the salient people (who might be the speakers, the audience, or the agents at issue). The pragmatic upshot of this is *quasi-expressivism*: unrelativized use of end-relational language will typically function to express attitudes and to have recommendatory illocutionary force just as expressivists maintain, but via their pragmatics rather than their semantics or conventional meaning.

This quasi-expressivism about unrelativized normative sentences provides the basis for my solutions to the other metaethical puzzles in Chapters 6-8. In Chapter 6, I argue that the overriding salience of the ends that are most preferred enables the end-relational theory to explain our thought and talk about what is good or ought to be done *simpliciter* and *all-things-considered*, in cases where multiple ends must be weighed against each other.

Chapter 7 addresses two separate objections to a broadly instrumentalist metaethical theory. The first problem is the status of moral utterances as “categorical imperatives”. For an end-relational theory, this is a challenge to explain why uttering implicitly end-relational sentences has a prescriptive function if the agent in question doesn’t care about the implicit end. I argue that here the salient ends are generally those of concern rather to the speaker (in paradigmatically engaged moral speech), and that this practice employs a rhetorical device of moralism whereby the speaker pragmatically expresses their own preferences in a prescriptive way.

Chapter 7’s second puzzle is to make sense, on an end-relational semantics, of our thought and talk about final value, or what is “good for its own sake”. On the basis of linguistic and psychological evidence I argue that by its semantics, ‘good for its own sake’ means *good for itself*—i.e. raises its own probability—and that by restriction of the relevant ends to those that are saliently preferred, the notion of final value amounts to what is relevantly an object of intrinsic desire.

Finally, Chapter 8 addresses the classic problem of lost disagreement for relativistic theories like COT’s, which imply that apparently disagreeing speakers actually assert compatible propositions (addressed to different ends). Drawing on the quasi-expressivist results of earlier chapters, I show that the end-relational theory predicts that these utterances will pragmatically express incompatible preferences,

recommendations, or prescriptions, providing a “disagreement in attitude” of the kind highlighted by expressivists. This informs a secondary sense of “confusion of tongues”: that fundamental normative disagreement is not over a single elusive set of truths, as moral philosophy often seems to assume, but involves a kind of talking past each other.

COT concludes with an overview of its implications for normative language, metaphysics, psychology, epistemology, ethics, and metaethics itself. To contextualize my critics: Dowell raises issues primarily about COT’s analytic methodology; Worsnip raises issues primarily about methodology, implications for normative ethics, finality (Ch. 7), and disagreement (Ch. 8); Koehn raises issues primarily about COT’s semantics for ‘good’ (Ch. 2), its treatment of conflicting ends (Ch. 6), and finality (Ch. 7).

### **Reply**

I am grateful to Alex Worsnip, Janice Dowell, and Glen Koehn for their challenging comments on my book *Confusion of Tongues* (OUP 2014).<sup>1</sup> As they side with me on several significant metaethical disputes—in particular we all accept a unifying contextualist semantics for most normative language—our disagreements may seem internecine to some readers, but there are many important issues remaining. In this reply I first address Worsnip’s case for contextualism without relativism. Next I address Dowell’s and Worsnip’s skepticism about whether COT succeeds in providing an analytic reduction of the normative, and Dowell’s recommendation to pursue an alternative, synthetic method. I then consider Worsnip’s comments on COT’s implications for normative ethical theory, and end by responding to Koehn’s challenges to the details of my end-relational semantics and treatment of the notion of final value.

### *Contextualism All the Way Down*

The dispute between relativism and absolutism (non-relativism) is fundamental in metaethics. Adopting Worsnip’s definitions:

*Relativism*: “There are no normative facts about which standards are ultimately correct, or which ends are the ones we ultimately ought to have...not merely by the lights of some further positive standard or end, but in some standard- or end-independent way.”

Perhaps the most ambitious project of COT is arguing for relativism from a contextualist semantics for normative language. Paraphrasing Worsnip:

*Contextualism*: The semantic content of sentences involving normative terms (like ‘ought’) varies depending on the context in which they are uttered, not because of

---

<sup>1</sup> I also thank Charles Côté-Bouchard, David Barnett, and especially David Clark for helpful comments, Tony Ellis and *Analysis Reviews* for providing this forum, and the ACU Office of the DVCR and the USC Provost’s Office for grant support.

ambiguity in the meanings of normative words, but because they have a single character that includes one or more parameters fixed by context.

I seek to achieve this unity by attributing to normative terms like 'ought' a semantic value involving a relation to at least one parameter, which can be either explicit or implicit. To say or judge that S ought to do A is always to say or judge for some relevant end *e*, that *in order that e*, S ought to do A. If normative sentences are always relativized to ends, doesn't relativism follow? Worsnip denies this, objecting that I overstate the metaethical significance of contextualism about normative words.<sup>2</sup> I'll now attempt to clarify why and to what degree I take end-relational contextualism to have this significance.

How can a contextualist make sense of judgments about what we "ultimately ought" to do? If this 'ought' has an end-independent sense, then we cannot account for it with a uniformly end-relational semantics. We could postulate a brute ambiguity in 'ought' between end-relative and end-independent senses, but as an ambitious contextualist, Worsnip shares my aversion to that disunifying answer. Rather, he suggests that end-relational contextualism doesn't entail relativism because the relevant ends for some normative judgments can be picked out in a non-relativistic way. We can make sense of thoughts about what we ultimately ought to do as concerning what we ought to do relative to the *objectively correct* ends (de dicto).

Here I could complain about a bait-and-switch: Worsnip promised an end-independent sense of 'ought', but provides instead an account on which 'ought' is always end-relative, but some ends are privileged with a special, nonrelative normative property. This would be nit-picky, however; I agree that an "objective" or nonrelative kind of correctness might provide a nonrelative kind of normativity. But now we really need to turn our attention to the normative word 'correct'. Worsnip's proposal requires attributing to it a nonrelative sense. If it has such a sense, then it is incorrect to suggest that "all normative claims can be understood under a single contextualist semantics". At best, all 'ought' claims can be given a single contextualist semantics, but we will need to recognize nonrelative senses for at least one other normative term.

As a contextualist about *all* normative terms, I think 'correct' itself needs an end-relational semantics, and I find the idea of a nonrelative sense of 'correct' mysterious. Wittgenstein famously made a similar observation almost a century ago:

The right road is the road which leads to an arbitrarily predetermined end and it is quite clear to us all that there is no sense in talking about the right road apart from such a predetermined goal. Now let us see what we could possibly mean by the expression, 'the absolutely right road'... The first thing I have to say is, that [it] is nonsense! (1997: 67-9)

To be an "ecumenical" contextualist of Worsnip's kind requires making an exception somewhere, and recognizing some normative terms (whether actual or merely possible) as

---

<sup>2</sup> This kind of response is also pressed in Laskowski 2014, Fogal 2016, Perl 2017, Worsnip 2018.

different from the rest by having nonrelative senses, whether ambiguously or nonambiguously. We might alternatively label it *half-hearted* or *compromised* contextualism.

Half-hearted contextualists owe us an answer to a difficult question: which normative terms should we be (exclusively) contextualists about, and which should we not? Worsnip favors 'correct', other possibilities include 'important', 'rational', and 'reason'. But this choice seems arbitrary. For example, Judy Thomson (2008) makes the opposite choice, taking 'correct' to be merely standard-relative and 'ought' to be standard-independent. COT makes a case that we should rather be contextualists *all the way down*, arguing that we have no good reason to treat any of these normative terms as significantly different in this respect than any of the others. If we are (end-relational) contextualists all the way down then, I believe, we are committed to relativism.

### *Analytic Reduction of the Normative?*

If uncompromising contextualism entails relativism, then so much the worse for uncompromising contextualism, one might think. None of my present critics make a case against *metaphysical* relativism about normative facts and properties on behalf of absolutist realism, but Worsnip and Dowell both object to COT's *conceptual/semantic* relativism about normative concepts and language. They complain that it is implausibly restrictive because it doesn't allow people to say or think the nonrelativistic things they clearly sometimes mean to say or think (even if those things are confused or fantastical). To insist on semantic uniformity at this cost would be to fetishize it indefensibly.

Dowell additionally objects on these grounds to my claim to have done what is generally agreed to be the impossible: to give an analytic reduction of normative facts and properties to complexes of nonnormative (or "natural") facts and properties. The end-relational facts and properties I identify as normative do not seem to her normative in any sense she or normative realists like Tim Scanlon and Derek Parfit would recognize. Worsnip suggests someone with my metaphysical skepticism about nonrelative normative properties would be better advised to adopt an error theory about some normative thought and talk, while Dowell constructively suggests rather that the end-relationalist abandon any claim to be doing conceptual or semantic analysis (the analytic method), and instead defend an end-relational reduction as a *synthetic* claim. She suggests this would strengthen the end-relational theorist's position, because my professed analytic method is both "independently implausible" and "inessential".

I'll start with a couple of concessions. First, my claims in COT to have reduced normative facts and properties are based on an assumption about what it means to call something "normative" that I do not now wish to vigorously defend.<sup>3</sup> This is the assumption that a fact or property is "normative" (in a "minimal" sense) if it is a fact or property of the kind that our normative

---

<sup>3</sup> See Finlay 2019; this is also conceded in COT on pp. 249-250

judgments are about. Finding that normative judgments are about end-relational facts and properties, I concluded that “normative facts and properties” can be reductively analyzed in end-relational terms. This implies that both of the following claims/judgments are at least minimally normative, being about the same kind of end-relational state of affairs:

- (1) If the lake is going to freeze, the temperature has to drop.
- (2) If we are going to survive, we have to get out of the lake.

I granted that (2) may be normative in a more “robust” or “full-blooded” sense than (1), which I explained in quasi-expressivist terms of its contingently action-guiding function and pragmatic implication of favorable attitudes, in virtue of its relationship to our concerns (2010; 2014: 79).

Plausibly, however, what prompts many philosophers to classify a sentence like (2) as normative is actually exhausted by whatever relevantly distinguishes (2) from (1), which in COT is the business of the pragmatic, quasi-expressivist side of my account.<sup>4</sup> In this case, what makes a judgment “normative” is not (merely) what it is about, but in part its being made from a motivationally engaged perspective. The right conclusion to draw would then perhaps not be that normative facts and properties can be reduced, but rather that *there are no normative facts and properties*—and, perhaps, no words that are normative by their semantics.<sup>5</sup> This would be despite our normative claims being both semantically descriptive (contra expressivism) and sometimes true (contra error theory). Since I think the adjective ‘normative’ can be and has been used to mean different things, I leave it to the reader to decide how best to describe COT’s metaphysical implications here.<sup>6</sup>

Second, I do not definitively dismiss the possibility that some people sometimes use “normative” language with a nonrelative meaning—uses for which I would accept an error theory as Worsnip suggests.<sup>7</sup> COT’s main dialectical thrust is not that the semantics and pragmatics of this language *proves* that our normative concepts are (ubiquitously) relativist. Rather, it is that the standard, influential objections to a relativist interpretation (e.g. from categoricity, finality, and disagreement) fail, because our first-order practices are substantially compatible with relativism. Even if one continues to doubt that *everyone’s* normative concepts could be relativistic, it significantly shifts the dialectic to show that we don’t have to be absolutists or error theorists about all or even most ordinary normative or moral judgments.

More stubbornly, however, I challenge the alleged evidence that ordinary folk and even absolutism-espousing philosophers really do have nonrelative normative concepts. I believe that those who (like Dowell and Worsnip) find this skepticism implausible aren’t being sufficiently sensitive to a crucial distinction between (i) our first-order normative concepts, or

---

<sup>4</sup> See Foster ms for a case that the difference is properly located in the semantics.

<sup>5</sup> Eklund 2016: 67, 71. The resulting view on normative judgment is most akin to the *relational expressivism* of Toppinen (2013) and Schroeder (2013).

<sup>6</sup> See Finlay 2019 for discussion. Another option is to say that while facts and properties are never normative *simpliciter*, they can be normative relative to particular motivated perspectives.

<sup>7</sup> On these issues see also my replies to Richard Joyce in Finlay 2008, 2011.

the meanings with which we use and interpret a normative term, and (ii) our second-order reflective understanding (theory, conception) of what we are doing or thinking when we use and interpret the term.

To deny that some people have absolutist theories of their first-order normative concepts, words, and practices would indeed be absurd. But since it is highly plausible that these concepts are somewhat opaque to our reflective understanding, it is not similarly absurd to deny that they are absolutist, at least once the compatibility of our practices with relativism has been demonstrated, as COT attempts. Dowell writes that “If our conception of the normative includes the realist’s, we will be forced to cast Finlay’s as an eliminativist reduction [since] end-relational facts do not accord with the realist’s conception.” My response is that while an end-relational theory may not accord with the absolutist realist’s reflective understanding (“conception”) of her practice, it may nonetheless be a correct account of even her concepts, or what even she means when she uses words like ‘ought’ and ‘good’. I believe this distinction largely explains the two different Finlays that Worsnip colorfully contrasts: the “careful conservative” wants to preserve as much of the first-order normative practice as possible. The “swashbuckling radical” seeks to expose common reflective or metanormative confusions about this practice, among both philosophers and the “folk”—which is the main point of the book’s title.<sup>8</sup>

Here I am telling many people that they are wrong about their own concepts and practices. Isn’t this preposterously arrogant? At the very least it’s a claim that needs strong support—which COT attempts to supply. One part of the case is charity: on the assumption that there is no nonrelative normative reality, the end-relational interpretation rescues the truth of many first-order normative claims from the threat of error theory (more on this below). But charity has its limits, and I don’t consider ordinary thought or practice immune from error—even systematic error. In COT and elsewhere I offer arguments and empirical evidence that ordinary normative concepts are end-relational/ relativistic. Here I’ll gesture briefly at some of this evidence.

The master-argument of COT is an inference to the best explanation: whole-hearted contextualism of an end-relational variety enables us to give complete, naturalistically impeccable explanatory solutions to a whole range of central metaethical puzzles. What are normative properties and how are they connected to the natural world (the “location problem”)? How do we detect them? How and why do they move us, and seem to have authority over us? Half-hearted contextualism (or absolutism) does not as easily provide answers to any of these puzzles.

---

<sup>8</sup> Unfortunately there is no space here to discuss complications involving possible interactions between practice and theory. But in response to Worsnip, note that (confusingly!) the secondary allegory of ‘confusion of tongues’, as our disagreements involving a kind of talking past each other, was not meant to imply that in doing so we ourselves are actually confused (as in the primary, metaethical allegory). Here Dowell’s distinction between a confusion *of* tongues (secondary allegory) versus a confusion *over* tongues (primary allegory) is very helpful.

As part of this explanation I've argued that what it is for a concept to have an intuitively normative character can be reductively analyzed in end-relational terms, as evidenced by the way that our sense for normativity in modal verbs like 'ought' is controlled by features of grammatical tense and aspect that syntactically function to represent a relation to an end (see Finlay 2010, 2014: 61f).<sup>9</sup> Compare:

- (3) Before you receive your license, you have to have taken the test.
- (4) Since you received your license, you have to have taken the test.

Here, the modal verb in the consequent of (3) is naturally intuited as normative (=required), while exactly the same modal verb in the consequent of (4) is naturally intuited rather as epistemic (=evidently), with the only difference between the sentences being that the antecedent of (3) temporally represents a state of affairs as in the relative future (an "end").

If this analysis is correct, then the very idea of a nonrelative normative concept is deeply incoherent, like the idea of a triangle with less than three angles. By contrast, the half-hearted contextualist faces a difficult question here. What do end-relational concepts and absolutist concepts have in common, that would lead us to classify both as "normative"? More generally, it seems to me that absolutists have no satisfying story to tell about where the sense for normativity comes from (to get a little polemical, they are left with the spooky *pricking of their thumbs*).

COT also offers empirical evidence in the form of our shared linguistic behavior that our normative concepts—even the concepts of absolutists—are thoroughly end-relational. A particularly striking example is the finding that while the sentence 'It is good for its own sake' is used quite often (4.6 million hits on a 2013 Google search, 390,000 in 2019), virtually nobody—*not even metaethical absolutists*—ever uses the string of words 'It is bad for its own sake' (one hit in 2013, six in 2019 including two mentions referencing COT). This is an otherwise baffling gap in usage that my end-relational interpretation of our shared normative concepts explains and indeed predicts: in brief, it implies for something to be "bad for its own sake" would be for it to lower its own probability, which nothing can do.<sup>10</sup>

More generally, assuming there are not actually any nonrelative normative properties in the world to which we could be sensitive, an explanation is needed for the fact that even people with absolutist theories still have numerous first-order normative intuitions and judgments. Merely believing in the existence of fairies isn't sufficient to explain particular perceptions of fairies. I argue that the end-relational theory provides the best explanation of everyone's normative intuitions and judgments, and is therefore plausibly correct about everyone's normative concepts—even those of avowed metaethical absolutists.<sup>11</sup> Analogously, if someone

---

<sup>9</sup> As conceded to Dowell above, some pragmatic seasoning may also be required.

<sup>10</sup> See COT pp. 197f for discussion and important qualifications; I hope to expand on and clarify these points in future work.

<sup>11</sup> Cf. Finlay 2008, 2011.



tells us that what they mean by ‘triangle’ is something other than being a regular polygon with three angles/sides, but we find them classifying as “triangles” all and only regular polygons with three angles/sides, then we should doubt their reflective understanding of their own pre-reflective concept of triangularity.

Notice that as described above, COT’s case for the end-relational theory turns almost entirely on what best explains *our ordinary first-order linguistic and mental behavior*. I believe metaethics is largely a quest for self-understanding, of what we (individually and collectively) are doing when we make normative judgments and claims. This is the profile of an *analytic* method, and of semantic and conceptual analysis—not of a synthetic method for investigating an independent and unknown metaphysical reality. If the evidence and arguments provided in COT are probative, this is excellent empirical evidence that at least one version of an analytic method is viable in metaethics. But if Dowell is right to dismiss the analytic method, then the evidence and arguments in COT must fail somehow, and an entirely different approach would be needed to vindicate the end-relational theory. The right way to adjudicate this disagreement, I believe, is to examine the individual arguments and pieces of evidence in COT on their own merits, rather than by deferring to general metaphilosophical or metasemantic considerations.<sup>12</sup>

While I can’t do justice here to Dowell’s recommended method of “vindicating deflationism”, my impression is that this is really a proposal to change the subject and come to use terms like ‘ought’ and ‘good’ in a new way, to refer to different properties and relations.<sup>13</sup> I believe this should only become a live option once we’ve justifiably decided there is something faulty about the way we presently use our terms, a decision we shouldn’t make before attempting conceptual analysis—and if COT’s analysis is correct, one that will then appear unmotivated.

### *Why Normative Ethical Theory is Bullshit*

Worsnip also raises an original challenge to my claim that COT’s relativistic contextualism is neutral with regard to normative ethical theory. Here, he suggests, my radical inclinations are in direct tension with my conservative inclinations. Drawing a useful distinction between (i) the truth of *first-order ethical claims* made by Utilitarians, Kantians, etc. and (ii) the truth of their *ethical theories* (Utilitarianism, Kantianism, etc.), he accepts that the view in COT is compatible with (i), but denies it is compatible with (ii)—or if it is, he suggests, it implausibly implies that *all* normative ethical theories are true, contrary to the shared assumptions of ethical theorists of all stripes—and therefore is anything but conservative about ethical theory. Further, he objects that the first-order implications of my reduction of ethical claims to empirical claims implausibly make ethical truth too cheap, having to concede to the truth not only of the contested moral

---

<sup>12</sup> I don’t have space here to explore general objections to the analytic method; see Finlay ms.

<sup>13</sup> Compare “reforming definition” (Brandt 1979, Railton 1986, Harman 1996), and “conceptual ethics” (Burgess & Plunkett 2013).

claims of opposed ethical theorists, but even of abhorrent and indefensible claims (see also Dowell).

COT does offer a reply to this objection, one which generalizes also to the concern about its normative non-neutrality. In passing, note that one solution compatible with the theory in COT appeals to the qualifier 'morally'. Ethical theory is generally concerned with what we *morally* ought to do. So one response is to argue that morality is objectively defined by a special set of ends.<sup>14</sup> However, plausibly the qualifier 'morally' is itself used as a normative term. (Perhaps the moral ends are just the *most important* ends, as Dowell has previously proposed.) Being a contextualist all the way down, I ultimately prefer to understand this relativistically, as Worsnip assumes: proponents of rival ethical theories just prefer different final ends. So I favor a different reply.

The heavy lifting in Worsnip's objections is done by the notion of theories or claims being *true*. They depend on interpreting the truth predicate in a purely cognitivist or correspondence way: a moral claim is "true" iff the descriptive proposition that was thereby asserted is true. But although COT describes its semantics as "descriptivist" and "cognitivist", this is an interpretation of 'true' I am committed to rejecting for ordinary normative discourse.<sup>15</sup> According to the *quasi-expressivism* advanced in COT, to make a (full-blooded/ engaged) normative claim is not merely to assert an end-relational proposition, but to do this from the perspective of a concern for the relevant end, and accordingly to express a conative stance or attitude.

Given familiar deflationist observations about the usual function of 'true' in ordinary discourse as a device of expressing endorsement, we can therefore distinguish between being true\* (truth\*) in a broad, metaphysically neutral and deflationist sense, and being true# (truth#) in a technical, narrowly correspondence sense. I don't dispute Worsnip's characterization of the implications of COT in the technical sense: the descriptive propositions asserted by competing ethical theorists could all be true#, depending on how we interpret their implicit relativization. But I maintain that in both ordinary normative discourse and normative ethical theory, the truth-predicate is standardly used and interpreted in the deflationary, engaged way, as true\*. This reading is necessary, for example, for the accusation that my theory has abhorrent and indefensible normative implications. There is nothing *morally* abhorrent or indefensible about asserting that using chemical weapons was the (instrumentally) best means to Saddam Hussein's preferred ends.

Where does this leave normative ethical theory? Assuming we are not amoralists, in judging whether a moral theory is "true" we are not (merely) judging the truth# of some empirical

---

<sup>14</sup> Different theories would then be disagreeing over what those ends are. We could then say that some ethical theorists are factually mistaken about what the moral ends are.

<sup>15</sup> See Björnsson & Finlay 2010: 30f, Finlay 2014: 236f. One potential source of confusion here is the ambiguity in 'semantics'. Read as about *truth-conditions*, my position may seem incoherent, but I rather mean *conventional signification*.

claim, but crucially (also) making/expressing a commitment about what matters to us, and what we stand for. Insofar as the commitments expressed by the proponents of rival ethical theories are incompatible, we cannot coherently judge them all to be “true” (true\*). We cannot coherently express an overriding concern *both* for maximizing happiness *and* for complying with a deity’s will, unless we happen to think these are corealisable. So I stand by the claim that COT’s relativist version of contextualism is compatible with accepting any of a wide number of ethical theories (as true\*), although not with accepting them all.

However, there is a way in which the view in COT will seem radically skeptical about first-order ethical theory to many philosophers. Worsnip contrasts the “enterprise of first-order ethics” as a philosophical debate with the “exercise of political rhetoric”, and questions what the point of the former could be on COT’s view. If we understand philosophy as the search for objective, perspective-independent truth (truth#), then I say that (at least fundamental) normative theory is either confused about its own central concepts, or isn’t really philosophy at all but more akin to political rhetoric: *bullshit* in the technical, nonpejorative sense defined by Harry Frankfurt, as being aimed at persuasion without concern for truth#. <sup>16</sup> But (echoing Wittgenstein) I don’t deny it can be *important* bullshit. Different understandings of philosophy itself are also possible, like Nietzsche’s concept of philosophy-as-legislation. COT additionally makes room for significant philosophical issues and debate over the consistency, implications, and empirical adequacy of the claims made by various ethical systems, and I am open to the possibility that carefully thinking these problems through under ideal conditions would lead the majority of people to converge in what ethical theories they accept as true\* in virtue of our common human nature (roughly as hypothesized, in different ways, by David Hume and Michael Smith)—though I don’t see much reason for confidence.

### *What Kinds of Ends, What Kinds of Relations?*

A common feature of work in metaethics that I find regrettable is speculation about big picture issues (e.g. is value real, how do we know about it, how does it motivate us, how is it related to the natural world) without first trying to nail down exactly what we mean/are talking about. Accordingly, COT starts by seeking a fine-grained semantics for our normative terms before turning to address metaethical puzzles. I am very pleased that Glen Koehn focuses critical attention on the details of my semantics.

Why do I call my theory ‘end-relational’—especially when there are other terms in the vicinity, such as ‘instrumental’, ‘teleological’, and ‘Humean’—and why ‘relational’ rather than simply ‘relative’? ‘End’ is here used as a technical term for the *content* of an actual or possible desire, intention, or similar attitude—or alternatively, a potential *outcome*. It contrasts with the desire or attitude itself, differentiating the view from standard Humean theories (and one prominent

---

<sup>16</sup> Frankfurt 2005; I take the suggestion from Mike Ridge. This stance is shared with noncognitivists going back at least to A. J. Ayer and Charles Stevenson.

reading of 'instrumental'), which locate normative facts in relations to our desires or other psychological attitudes, making it easier to accommodate epistemic and moral value, for example. This doesn't yet draw a distinction from 'teleological'; however that term is closely associated with the Aristotelian idea that nature itself sets proper ends for things (as Koehn discusses), so a different term is useful at least insofar as it detaches that connotation.

Why "relational"? This is chosen to highlight an ontological claim absent from the vaguer terms 'relativ(ism)' 'teleological', and 'instrumental'. The claim is that normative facts, properties, etc. consist in *relational properties*, i.e. properties of *standing in some relation* to some other thing. However, I agree with Koehn that goodness is not a relation (a 2-or-more-place operator, like *better than*). Rather, it is a relational property (a one-place operator). Or more correctly, it is *many* relational properties, because goodness always consists in a relation to an end, but there are countless different ends (goodness-for-e1, goodness-for-e2, etc.).<sup>17</sup>

Koehn raises two main objections to my version of an end-relational theory. He challenges (i) my account of ends as particular states of affairs or outcomes, and (ii) my view that the essential relation involved—at least in the cases of 'good' and 'ought'—is that of *raising probability* (of the end). I'll address these in turn.

When we talk about goodness, Koehn writes, we are often not talking about individual ends construed as particular states of affairs, but rather "of anything of a certain type". So for example, to say that Maria is a good tennis player is not to talk about "a relation between Maria and any particular thing or event". This seems importantly right; what can I say to address it?

Note that Koehn focuses on one particular kind of construction involving 'good', its "attributive" use to ascribe goodness to particular people or things ('a good K'). We use 'good' in a variety of other grammatical constructions (e.g. 'is good for', 'it is good that'), and in Chapter 2 COT seeks an ambitiously unifying syntax and semantics underlying all these different forms. I argue that 'a good K' requires considerable syntactic transformation to find its underlying form, taking as more basic a form in which 'good' takes propositional arguments: 'It is good for e that p.'<sup>18</sup> As Koehn observes, to be a good K is to be good *for* or *at* something, and we interpret 'good tennis player' (good at playing tennis) in terms of winning outcomes: the generic goal of a person playing tennis is to defeat their opponent. I therefore suggest an interpretation along roughly the following lines:

(5) It is good for *tennis opponents' being defeated*, that *Maria is the person playing tennis against them*.

---

<sup>17</sup> 'Good for' does express a relation, however, between at least its object and an end.

<sup>18</sup> This requires tolerating quite extensive syntactic transformations. If one has qualms, these claims can be reinterpreted as identifying the connections between slightly different senses (polysemies) of 'good'.

The end-term here, ‘tennis opponents’ being defeated’, is a nominative phrase apparently referring to a kind of event or outcome. We could understand it with a generic quantifier, roughly as follows:

(5') For typical opponents  $x$  and typical tennis games  $y$ , Maria playing in  $y$  against  $x$  is good for  $x$  being defeated in  $y$ .

This sentence is true, on my probabilistic semantics, in case the outcomes in which typical opponents are defeated are more likely if Maria is the person playing against them than otherwise. This seems to me a promising, albeit approximate analysis. Compare how sports statisticians measure a player’s quality in terms of “Wins Above Replacement” (WAR): how many more wins are expected by substituting the player for another who is league-average. So while I agree with Koehn that generic evaluations like ‘good at tennis’ can’t be analyzed by single ends or outcomes, I am optimistic they can be analyzed in more general end-relational terms involving quantification over the end-parameter.

Second, Koehn challenges my particular account of the relation involved in terms of probability-raising. One question is whether the relevant probabilities are subjective/epistemic, or objective/metaphysical (“chances”). My answer is *Both*. It’s widely accepted that we have both subjective and objective senses of ‘ought’ and reasons, which in my view corresponds roughly to which sense of probability we have in mind. To avert worries about this introducing deep disunity in our normative concepts, COT (pp. 42-5) sketches an ambitious reduction of probability to a measure of a space of possibilities that can be defined variously in subjective ways (by appeal to some subject’s information) and in objective ways (by appeal to some set of facts, such as the state of the world at some time). One implication is that construing normative properties probabilistically does not make them essentially characterizable in doxastic or psychological terms: probability as such does not involve psychological attitudes such as confidence.

At the heart of Koehn’s disagreement with my probabilistic semantics is his observation that evaluative judgment is not always concerned with probabilities, as when we say ‘Good dog!’ In its place, he proposes a rival end-relational theory that “to say something is good is to say that it is such as to realize, or has the capacity to realize, a given end,” a view that “takes power, capability, or sufficiency, rather than probability, as fundamental to goodness”.

I concede that my semantics is a bit strained in application to uses like ‘Good dog’ or ‘good breakfast’. The adjective ‘good’ is used in a very wide variety of ways, and so probably any unifying analysis is going to “strain its concept” in some cases (Gibbard 1990). Of course, that might be reasonably be taken as a strike against a unifying analysis, but my strategy in COT is to see how far one can take us. My hypothesis is that these uses are outliers that the theory can easily make sense of. I take sentences of the form ‘It is good for  $e$  that  $p$ ’ as central, and for these I argue that the probability-raising account is the best fit. This choice of focal case may be suspected arbitrary. However, these are the uses with more syntactic structure (compared with

‘Good dog’, for example), and so are plausibly the most revealing of the precise meaning of ‘good’. By contrast, it seems hopeless to begin with an analysis tightly tailored to ‘Good dog’, and then try to extend this to make sense of more complex sentences like ‘It is good for  $e$  that  $p$ ’ (the traditional noncognitivist direction of analysis).

Here’s a conjecture about how ‘good dog’ and ‘good breakfast’ can be assimilated on a probabilistic semantics. Given that ‘ $p$  is good for  $e$ ’ means (per COT) that  $p$  raises the probability that  $e$ , it is unsurprising that ‘good’ would occupy its role as our most general term of endorsement. If you want  $e$ , then you’ll naturally be looking for things that are good for  $e$  in this sense. So things called “good” are normally things to be sought, chosen, preferred, etc. I suggest this is sufficient to make it nonmysterious why we would say things like ‘good dog’ and ‘good breakfast’ even when we’re not thinking about probabilities.

By contrast, Koehn’s rival theory involving power or sufficiency for an end does not seem so plausibly extended to cover all the cases.<sup>19</sup> Something can be good for an end without being sufficient for it, or having the power to bring it about. For example, scoring runs in the first inning of a baseball game seems clearly good for winning the game, but isn’t sufficient for that end nor does it have the power of bringing it about. By contrast, this goodness seems naturally understood in terms of probability-raising. On the other hand, the probability-raising analysis easily subsumes the sufficiency cases. If  $p$  is sufficient to make it the case that  $e$ , then  $p$  raises the probability that  $e$  (against a suitable background).<sup>20</sup>

Koehn acknowledges that some uses of ‘good’ do appear to fit a probabilistic analysis, but proposes that in these cases, probability is correctly located in the *ends* rather than in the relation to the end: that if we say that scoring runs early in a baseball game is good for winning, we can interpret this goodness not in terms of raising the probability of the end of *winning*, but rather in terms of sufficiency or power for the end of *raising the probability of winning*. These ends are rationally connected: if you have the end of winning a baseball game, then you will rationally also desire the end of raising the probability that you win, and therefore of scoring runs in the first inning.

Even conceding this, however, it doesn’t accommodate the natural thought that scoring runs in the first inning is good *for winning the game*. Another problem is that any *rational* connection between desiring an end and desiring some promotive means itself plausibly needs a normative explanation. *Why* should you rationally care about scoring runs in the first inning, if this isn’t good relative to any of the ends you already do (or should) care about? On the probability-

---

<sup>19</sup> I favored such a view (in terms of *satisfying an interest*) in my 2001, but abandoned it for this reason.

<sup>20</sup> There is a major issue here, corresponding to the debate between evidential and causal decision theory, over whether goodness also requires a causal connection. See Behrens & DiPaolo 2011 for a related objection to my analysis of “promoting an end”. I allow that a causal connection may be necessary for some *kinds* of goodness, but also accept the existence of non-causal goodness; see COT p. 43.

raising semantics, the goodness of the means more intuitively precedes and helps explain the rational motivation for it.

A further motivation for a probabilistic analysis is the *gradability* of goodness, as in degrees of value, or *betterness*. Degrees of goodness for an end are plausibly degrees of probability-raising, but power or sufficiency are not so obviously amenable to gradability. In response, Koehn correctly points out that this probabilistic treatment of betterness faces a serious problem in cases where there are competing ends: “The better coffee is such as to bring about [a more preferred] end, rather than being held to increase the chances of the same end being realized... Thus, ‘better than’ is at least sometimes to be understood as revealing sufficiency for a preferred end rather than revealing a greater chance of realizing the same end.” I agree that ‘better’ sometimes involves a more preferred end, (apparently) rather than greater probability of the same end. But here Koehn’s theory and mine are on a par. If ‘good’ means *capable/sufficient of producing the end*, then ‘better’ can’t mean *more good*, since *more capable/sufficient of producing the end* isn’t a plausible analysis of these cases.

Koehn proposes a noncompositional solution: ‘better’ in these cases doesn’t signify a greater degree of the same relation, but rather a different relatum/end. But this move is also available to a probabilistic analysis: i.e. ‘better’ doesn’t in these cases signify raising the probability of the same end more, but of raising the probability of another, preferred end. COT (Chapter 6) proposes this solution, and also attempts to recover compositionality by appealing to a principle (“Preferential Selection”, p. 152) that the salient end is generally the *contextually most preferred* end. So I propose that ‘Bean X is better than bean Y’ is used to say that *[using] bean X raises the probability of achieving the more preferred flavor more than [using] bean Y does*. In cases of gradability involving different ends, we may have a stalemate. However, in cases involving a single end, I think the probabilistic analysis is more successful. In summary, what leads me to embrace a probabilistic semantics is the belief that it best accommodates all the cases, while acknowledging there are outlying cases where appeal to probabilities seems gratuitous.

*Finally...*

Koehn’s final objection is to COT’s treatment (in Chapter 7) of final ends and goodness, or what is “good for its own sake”, which he reasonably identifies as the “main stumbling block” to accepting an end-relational theory (similar concerns are suggested by Worsnip and Dowell). My proposal, based in part on a compositional analysis of qualifier ‘for its own sake’, is that to say that something is “good for its own sake” is to say (in part) that it is good for itself—i.e. raises its own probability. To say that pleasure is good for its own sake is to say that (one’s having) pleasure raises the probability of (one’s having) pleasure. Koehn speaks for many, I’m sure, when he writes that my view makes no sense, although I think the problem doesn’t strictly involve intelligibility but rather plausibility as it appears to attribute final value to everything,

trivially. But as Koehn himself accepts an end-relational theory he does not invoke an absolute or nonrelational property of final goodness. Rather, to be “good for its own sake” is to be “good (for some end in view) without needing anything further.”

An initial point is that I want to distinguish between things being good for their own sakes, and things being good in themselves. Regarding “good in itself” I agree with Koehn: this means to be good (for some unspecified end) solely on the basis of the thing’s own intrinsic properties and independent of extrinsic relations to other things (see COT p. 200). Something can be good in itself (for some end) without being good for itself/ for its own sake.

I think the triviality objection only seems cogent if one again ignores the quasi-expressivist component of my analysis. According to this, engaged/full-blooded normative judgments do not merely assert end-relational propositions, but crucially also pragmatically express preferences and concerns. While I claim there is a trivial sense of ‘good for its own sake’ that applies to everything (“What is having belly button lint good for? Well, it’s good for having belly button lint. It’s good for its own sake,”) the cases of interest to philosophers are engaged judgments, of what is Good (relative to a saliently desired or preferred end). To say that something is Good, for its own sake, is according to COT to say that it raises the probability of some end *e*, where *e* is indicated to be *relevantly desired/preferred*, and to further clarify that the end *e* in question is just the thing itself. If it is worried that this makes the connection between actual desire and judgments of final value too tight, note that “relevantly desired” can be interpreted in counterfactual or idealized ways (e.g. desired under conditions of full information), which allows for our talk about final good to be about that which is *desirable* or *fit for desire* rather than simply what is desired. However, I deny that final value gives us something to care about or aim at; rather, as living creatures we already inescapably have various things that we care about or aim at, and this is why we are beings who engage in normative thought at all.

In COT I maintain that this is not merely an adequate analysis of ‘good for its own sake’, but also that there is compelling empirical evidence in its support. In addition to grammatical evidence from compositionality (of the construction ‘for its own sake’), moral psychology (the near-platitude that things judged good for their own sakes are things desired for their own sakes), there is the previously mentioned, striking pattern that people have used the sentence ‘It is good for its own sake’ often while ‘It is bad for its own sake’ effectively never—a peculiarity which my theory predicts but otherwise so far lacks an adequate explanation.

Here as elsewhere, the theory advanced in COT may seem to many readers implausible and even absurd as an account of their own concepts, but we should accept it nonetheless so long as it predicts and explains our ordinary thought, talk, and behavior better than any rival.



## References

- Behrends, J. & DiPaolo, J. (2011). "Finlay and Schroeder on Promoting a Desire." *Journal of Ethics and Social Philosophy* 1: 1-7.
- Brandt, R. (1979). *A Theory of the Good and the Right*. Oxford University Press.
- Burgess, A. & Plunkett, D. (2013). "Conceptual Ethics I." *Philosophy Compass* 8(12): 1091-1101.
- Eklund, M. (2016). *Choosing Normative Concepts*. Oxford University Press.
- Finlay, S. (2001). *What Does Value Matter? The Interest-Relational Theory of the Semantics and Metaphysics of Value*. University of Illinois PhD dissertation.
- Finlay, S. (2008). "The Error in the Error Theory." *Australasian Journal of Philosophy* 86(3): 347-69.
- Finlay, S. (2011). "Errors Upon Errors: A Reply to Joyce." *Australasian Journal of Philosophy* 89(3): 535-47.
- Finlay, S. (2010). "Normativity, Necessity and Tense: A Recipe for Home-baked Normativity," in R. Shafer-Landau (ed.), *Oxford Studies in Metaethics* 5. Oxford University Press, pp. 57-85.
- Finlay, S. (2014). *Confusion of Tongues: A Theory of Normative Language*. Oxford University Press.
- Finlay, S. (2019). "Defining Normativity," in D. Plunkett, K. Toh, & S. Shapiro (eds.), *Dimensions of Normativity*. Oxford University Press.
- Finlay, S. (ms). "The Empirical Armchair."
- Fogal, D. (2016). Review of *Confusion of Tongues*. *Ethics* 127(1): 281-8.
- Foster, J. (ms). "Not-At-Issue: The Normative 'Flavor' of Instrumental Necessities."
- Frankfurt, H. (2005). *On Bullshit*. Princeton University Press.
- Harman, G. (1996). "Moral Relativism," in G. Harman & J. J. Thomson, *Moral Relativism and Moral Objectivity*. Blackwell, pp. 1-64.
- Laskowski, N. (2014). "How to Pull a Metaphysical Rabbit out of an End-Relational Semantic Hat," *Res Philosophica* 91(4): 589-607.
- Perl, C. (2017). *Positivist Realism*. USC Ph.D. dissertation.
- Railton, P. (1986). "Moral Realism." *Philosophical Review* 95(2): 163-207.
- Schroeder, M. (2013). "Tempered Expressivism," in R. Shafer-Landau (ed.), *Oxford Studies in Metaethics* 8. Oxford University Press, pp. 283-314.

- Thomson, J. J. (2008). *Normativity*. Open Court.
- Toppinen, T. (2013). "Believing in Expressivism," in R. Shafer-Landau (ed.), *Oxford Studies in Metaethics* 8. Oxford University Press, pp. 252-82.
- Wittgenstein, L. (1997). "Lecture on Ethics," reprinted in S. Darwall, A. Gibbard, & P. Railton, *Moral Discourse and Practice*. Oxford University Press.
- Worsnip, A. (2018). "'Ought'-Contextualism Beyond the Parochial." *Philosophical Studies* 176(11): 3099-3119.