

Oughts and ends

Stephen Finlay

Published online: 2 February 2008
© Springer Science+Business Media B.V. 2008

Abstract This paper advances a reductive semantics for ‘ought’ and a naturalistic theory of normativity. It gives a unified analysis of predictive, instrumental, and categorical uses of ‘ought’: the predictive ‘ought’ is basic, and is interpreted in terms of probability. Instrumental ‘oughts’ are analyzed as predictive ‘oughts’ occurring under an ‘in order that’ modifier (the end-relational theory). The theory is then extended to categorical uses of ‘ought’: it is argued that they are special rhetorical uses of the instrumental ‘ought’. Plausible conversational principles explain how this end-relational ‘ought’ can perform the expressive functions of the moral ‘ought’. The notion of an ‘ought-simpliciter’ is also discussed.

Keywords Ought · Normativity · End-relational theory · Semantics and pragmatics · Hypothetical imperative · Categorical imperative · Probability · Reductive naturalism · Deontic modals

As this *ought*, or *ought not*, expresses some new relation or affirmation, ‘tis necessary that it shou’d be observ’d and explain’d’ and at the same time that a reason should be given; for what seems altogether inconceivable, how this new relation can be a deduction from others, which are entirely different from it. (David Hume, *Treatise* III.i.1)

Hume never actually asserted what has come to be called ‘Hume’s Law’: that you cannot validly derive a proposition containing an ‘ought’ from propositions stating only what *is* the case. But he did pose a challenge that many have thought impossible to meet: to analyze ‘ought’ purely in terms that are not normative,

S. Finlay (✉)
School of Philosophy, University of Southern California,
3709 Trousdale Parkway, Los Angeles, CA 90089, USA
e-mail: finlay@usc.edu

evaluative, or deontic. It is not difficult to provide plausible analyses of ought-propositions into propositions that state what is the case. For example

Reasons-Ought: X ought to φ = (df.) X has more reason to φ than to act in any other way instead;

Value-Ought: X ought to φ = (df.) It is better that X φ s than that X does anything else instead.

These analyses, however, are not constituted wholly of terms that are ‘entirely different’ from ‘ought’: ‘reasons’ and ‘better’ are similarly normative, and similarly thought to be in need of analysis or explanation by those who see a need to analyze ‘ought’. Philosophers who follow G. E. Moore in rejecting ‘analytic naturalism’ about normative terms deny that any other kind of analytic definition is possible. In a recent statement of this view, Allan Gibbard writes that “no correct definition can break out of the normative circle, a circle of ought-like terms.” (2003, p. 6)

I shall propose a theory of the semantics of ‘ought’, consisting of six theses, which I call the *end-relational theory*. This theory is reductive or broadly naturalistic, decomposing ‘ought’ into a complex of nonnormative terms or concepts, and is thus a ‘cognitivist’ account, although it has a significant expressivist or noncognitivist element. There are many reasons for wanting such a reductive analysis, but I shall not address them here; the case for the theory will be based solely upon its intrinsic theoretical virtues and its plausibility and fit with the data. It claims to be interpretive rather than revisionary: to identify the meaning with which ordinary speakers use the word. It also claims to be exhaustive, capturing all legitimate uses of ‘ought’. This virtue, which distinguishes it from all competing accounts presently finding favour with philosophers, provides the principal argument given here in its support.

1 The instrumental ought

The method I propose is to begin by examining the occurrences of ‘ought’ that seem most amenable to reductive analysis, finding a plausible account of these, and then examining how much of the remaining data of the use of ‘ought’ this account can be extended to cover. It is fairly uncontroversial that the least problematic uses of ‘ought’ for the reductive semantic naturalist are instrumental uses, in ‘hypothetical imperatives’ like

- (1) If you want to live to see your grandchildren, you ought to stop smoking,
- (2) If Max is trying to evade arrest, he ought to mingle with the crowd.

This is not to say that instrumental oughts are unproblematic for the reductivist; recently there has been a flurry of philosophical activity aimed at how instrumental oughts are to be understood,¹ and some argue that the reductivist project fails even

¹ E.g., Korsgaard (1997), Broome (1999), Wallace (2001), Schroeder (2004, 2005), Setiya (2007), Dreier (2006), Finlay (forthcoming). Philosophers have traditionally viewed hypothetical ‘oughts’ as much less problematic than categorical ‘oughts’ (Kant, Wittgenstein and Mackie, at least, took this view), but many recent attempts to explain hypothetical ‘oughts’ have sought to reduce them to categorical ‘oughts’.

here. However this is the thin end of the wedge: almost everybody, reductivist and nonreductivist alike, agrees that there are such instrumental facts. To deny that sentences like (1) and (2) are used to express true propositions sounds like philosophical fancy run amok.

What is the proper analysis of hypothetical imperatives? On the surface, they are conditionals, the antecedents of which concern some agents' wants, desires, wishes, likes, intentions or plans, the consequents of which state that the agents ought to act in some way.² The need to state further conditions is evident from examples like

(3) If you desire to smoke a cigarette, you ought to chew gum.

In a genuine (and true) hypothetical imperative, the action that the consequent states an agent ought to perform stands in an instrumental relation to the end specified in the antecedent; the action must be one that *promotes* that end. Ceasing to smoke promotes its being the case, for some people, that they live to see their grandchildren, but your chewing gum does not promote its being the case that you smoke a cigarette. A further condition is that the action must be in some sense the *best* way, out of a relevant set of alternatives, to promote that end. If Max could possibly evade arrest either by mingling with the crowd or by slipping out the back door, and the latter option is the surer bet, then (2) is false: instrumentally, Max ought rather to slip out the back door.

This surface diagnosis of hypothetical imperatives is problematic. If they genuinely are conditionals, then presumably they allow 'detaching'. If Max is indeed trying to evade arrest, it follows simply from this and (2) that

(4) Max ought to mingle with the crowd.

But what if Max ought not to be trying to evade arrest in the first place? Perhaps he morally ought to surrender to the police, or is capable of quickly establishing his innocence. Surely it can't be the case that we ought to take the means to the ends that we ought not to be intending. This has the effect of legitimizing the pursuit of anything we might intend (etc.), no matter how foolish, wicked, or unwelcome: the poisoning of our children, gambling all our worldly possessions, etc.³

There are many solutions proposed in the literature. According to one, hypothetical imperatives can only be true if it is true that the agent ought to intend the end.⁴ According to another, they are to be understood as follows: Everyone ought to make it the case that (if they intend the end then they perform the means).⁵ You can bring yourself into compliance with this requirement in three ways: by performing the means, by ceasing to intend the end, or by acting to eliminate the efficacy of the means.⁶ According to yet another, the strain of truth in the hypothetical imperative is merely that if a person desires an end, then they have a

² E.g. von Fintel and Iatridou (2006a).

³ Another problem, which is a focus of concern for von Fintel and Iatridou (2006a), arises when a person has conflicting desires or goals.

⁴ Korsgaard (1997), Hampton (1998), Hill (1973).

⁵ Broome (1999), Hill (1973).

⁶ Schroeder (2004, p. 339).

reason to perform the means.⁷ These strategies are all revisionary: they all seek to weaken the claim in the imperative's consequent that we *ought* to perform the means—by qualifying it with conditions, or introducing other options, or diluting the normative stringency. The solution proposed in this paper leaves this claim intact, taking the consequent at its semantic face value.

Another common locution in which 'hypothetical imperatives' sometimes appear can be seen in

(5) In order to evade arrest, Max ought to mingle with the crowd.

Compare (2) and (5); it seems that (5) expresses a proposition that is either the same as or very similar to that which a speaker would mean to express by asserting (2).⁸ But (5) isn't even obviously a conditional. The 'in order to...' locution seems to qualify the 'ought' in some other way. And if (5) is not a conditional, then it won't necessarily permit detaching; that is, (5) and the fact that Max is trying to evade arrest may not entail (4), that Max ought to mingle with the crowd.⁹

It will be useful to reformulate this modifier so that it takes propositions as arguments. We might rewrite (5) as

(6) In order that Max evades arrest, it ought to be the case that he mingles with the crowd.

This is clumsier, but will allow us to handle more of the phenomena. It is widely disputed, however, whether (5) and (6) genuinely express the same proposition. Some contend that on some occasions of normative use, 'ought' expresses not a propositional operator, but a relation between agents and actions.¹⁰ I maintain that the propositional view can accommodate these uses, but here I must just assume it. In what follows, I shift between talking about ends, states of affairs and actions, and talking about the propositions in which they feature.

The first thesis of my end-relational theory of 'ought' is that all instrumental ought-propositions can be and are (more) transparently expressed in that form:

First Thesis: All instrumental ought-propositions can be transparently semantically expressed with an 'in order that...', modifier.

Ought-propositions with the modifier appear superficially to have the following features: (a) They apparently are not conditionals, and do not permit detaching the 'consequent', whatever the agent's desires or intentions might be. (b) The modifier seems to function to qualify the *sense* of the 'ought'; in (6) it is 'ought (in order that Max evades arrest)'. It is due to this that (6) is compatible with

⁷ Schroeder (2004), see also Mackie (1977, pp. 74–75).

⁸ See also Sæbø (2001) and discussion in von Fintel and Iatridou (2006a). An explanation would still be needed of why we say 'If you want...', if there are more perspicuous ways to express what we mean. One possibility is that it functions partly as a 'biscuit conditional' (If you want some biscuits, there's some on the table): i.e. not to mark the conditional truth of the consequent, but its conditional relevance. See also Dreier (2006); von Fintel and Iatridou (2006a).

⁹ See also Sloman (1970).

¹⁰ Harman (1973), Broome (1999), Geach (1991), Schroeder (2005, 2007). For the opposing view, see also Williams (1981), Wedgwood (2006).

- (7) It ought not to be the case that Max mingles with the crowd;
 (8) In order to be arrested, Max ought not to mingle with the crowd.

That is, ‘ought (in order that Max evades arrest)’, ‘ought (in order that Max is arrested)’, and ‘ought all-things-considered’—whatever that means—have different senses, which we can represent by subscripts: ought_{e_1} , ought_{e_2} , etc. (Features (a) and (b) are actually misleading. I shall argue that hypothetical ought sentences do indeed involve a certain kind of conditional, and that ‘ought’ itself only has one sense, while explaining why superficially it appears otherwise.)

To continue, (c) it is evident that ‘in order that e ’ is instrumental or teleological in function. It expresses a relation that takes as its e -argument propositions representing goals or potential states of affairs, and is approximately synonymous with ‘To make it the case that e ,’ and ‘If it is to be the case that e .’ We have observed that in genuine ‘hypothetical imperatives’ the counseled action stands to the end as some kind of means. We can conclude that each sentence of the form ‘In order that e it ought to be that p ’ says something about the instrumental significance of its being the case that p for its coming about that e .

So what is the semantic function of ‘ought’ within such sentential contexts? It is fruitful here to move sideways: ‘ought’ has complexities we can temporarily avoid by considering some of its fellow normative terms. As a modal auxiliary verb, ‘ought’ belongs in the same family as ‘must’, ‘have to’, ‘need to’, ‘may’, ‘might’, ‘can’, and ‘could’. It is interesting that all of these terms are both ordinary modals, denoting variously alethic (logical, metaphysical, nomic, circumstantial) and epistemic necessity and possibility, and also *normative*, deontic terms like ‘ought’; we talk in this regard of normative, deontic, or moral ‘necessity’ and ‘possibility’.

I will treat the (nonnormative) modals generically, as follows:

Modal Must: It must be the case that $p =$ (df.) Every possible world in which circumstances/laws C obtain is such that p , and

Modal May: It may be the case that $p =$ (df.) At least one possible world in which circumstances/laws C obtain is such that p .

Logical modality addresses compatibility (of representations) of worlds with the laws of logic alone; metaphysical modality addresses compatibility with metaphysical laws, nomic modality addresses compatibility with scientific laws, and epistemic modalities address compatibility with what is believed or known by some subjects, or knowable independent of other beliefs. Everyday use of modals is frequently circumstantial, addressing what is necessarily or possibly the case given certain circumstances. We might maintain, for example, that given electoral discontent with the present government, there *has to* be a change in the governing party in the next election. ‘Possible’ in these formulae signifies merely satisfaction of the criteria in, or being not ruled out by C ; what is epistemically possible need not be metaphysically or logically possible. There are therefore many different modalities. For simplicity I will contrast normative modals with ‘ordinary’ modals, meaning any and all of these other forms of modality.

Observe what happens when these modal terms appear under an ‘in order that ...’ modifier:

- (9) In order to earn money, you must have a job;
 (10) In order that I fulfill my duty, I have to sentence you to death;
 (11) In order to sell your home, you could try advertising it online;
 (12) In order that Sammy entertains his guests, he may not show home movies.

Are the terms used here as ordinary modals, or are they normative? The answer appears to be: they are *both*. (9) makes the alethic claim that every world compatible with the circumstance that you earn money (and other implicit conditions) is one in which you have a job. But we seem to have, in the same proposition, an (instrumental) normative ‘must’.¹¹ At least, the characteristic marks of normative necessity are present: the necessity of your having a job is compatible with your *not* actually having a job, and provided that the antecedent identifies a goal of the agent, the ‘must’ functions to guide and criticize behaviour rather than to describe (actual) behaviour. If we sought instrumental normative counterparts to these ordinary modal sentences (propositions), the results would differ in no distinguishable way from these. Hence:

Second Thesis: Normative modal terms (concepts) in their instrumental uses are simply ordinary modal terms (concepts) occurring under an ‘in order that...’ modifier.

In other words, instrumental normative terms are ordinary modals occurring in an *end-relational* context.¹² Instrumental normativity is reducible to the modal relation between some action and some end. To say that I ‘have to’ sentence you to death (in order to fulfill my duty) is to say, approximately, that if I don’t sentence you to death, then it is impossible (given my circumstances and the dictates of justice) that I fulfill my duty. To clarify and precisify, I suggest the following analysis. Supposing that *Modal Must* and *Modal May* are correct, then

Must_e: In order that *e* it must be the case that *p* = (df.) Every possible world in which circumstances/laws *C* obtain, *including its being the case that e*, is such that *p*;

May_e: In order that *e* it may be the case that *p* = (df.) At least one possible world in which circumstances/laws *C* obtain, *including its being the case that e*, is such that *p*.

‘In order that *e*’ affects the modal proposition here by introducing the end in *e* into the circumstances *C*, restricting the relevant worlds to those in which the end in *e* is realized. (*C* will incorporate other restrictions, including compatibility with natural laws and with the circumstances obtaining immediately prior to the event or action in *p*; what we must or may do depends on the circumstances in which we are to act).

¹¹ My experience has been that linguists and philosophers of language take this point as obvious (e.g., von Stechow and Iatridou (2006a)), and tend to query why it needs any argument. Philosophers working in metaethics, however, tend to find it startling and controversial.

¹² Kratzer (1977) offers a similar unifying semantics for ordinary and normative modals ‘must’ and ‘can’ that has been influential in linguistics. Her account of the normative terms is not reductive in the metaethical sense, however, as she retains unanalyzed concepts of duty and of goodness in her analyses of categorical and hypothetical imperatives, respectively.

It might be thought that ‘in order that *e*’ has a further significance, requiring that it is the case that *e* partially *because* it is the case that *p*. For otherwise instrumental requirements would include logical, mathematical, metaphysical and nomic necessities, as well as necessary consequences of the end, which seems odd. If the modifier does indeed have this significance, it would be easy enough to accommodate it by adding to the analysis that its being the case that *e* results from its being the case that *p*. This is unnecessary, however, because there are good reasons to think that the pragmatics of ordinary conversation provide an assumption of such a dependence.¹³ These troublesome instrumental claims would therefore be odd, but not false.

The end-relational theory provides a straightforward analysis of instrumental normative language, unifying the language of ordinary modality and normativity, and providing a univocal semantics for two isomorphic sets of terms. All else being equal, a theory that can unify multiple phenomena under a common principle is to be preferred over theories appealing to distinct principles, each of which is adequate to the phenomena in its own domain.

But what about ‘ought’? This normative term *par excellence* presents two difficulties. First, it may be thought that it anomalously lacks a nonnormative modal function. And second, it is difficult to see what modal relation ‘oughtness’ could be reducible to. But ‘ought’ *does* in fact have a nonnormative modal counterpart: the so-called *predictive* or *epistemic* ‘ought’, as we find in

(13) It ought to rain tomorrow

(14) Kitty ought to be arriving soon.

Significantly, ordinary modal and (unqualified) normative ought-propositions can be expressed by (virtually)¹⁴ all and only the same sentences. (13) is naturally read as an ordinary modal given that we seldom see fit to make assertions about how the weather ought normatively to be. But (14) is neutral between a normative and an ordinary modal reading. Based on theoretical symmetry, I propose:

Third Thesis: The normative ‘ought’ in instrumental uses is simply the ordinary modal ‘ought’ under an ‘in order that...’ modifier.

The proposal is that the normative and ordinary (‘predictive’) ‘oughts’ share a common semantics. This contradicts a widespread consensus among metaethicists that here we have two entirely distinct meanings for the same word.¹⁵ I shall first examine the function of the normative ‘ought’, then show how it enables us to derive a *prima facie* plausible account of the nonnormative ‘ought’.

¹³ I suggest below that these instrumental claims are a kind of conditional, after all. We are ordinarily disposed to reject as false or peculiar those conditionals in which the consequent is not dependent on the antecedent, such as ‘If Max is a cat, then two is the square root of four’. A standard explanation of this is as due to a conversational expectation of relevance, and in the present case we should observe that logical truths are generally irrelevant to the prospects of attaining some end.

¹⁴ von Fintel and Iatridou (2006b) argue that quantification is disallowed from taking wide scope over epistemic but not deontic ‘oughts’. This would support rather than undermine the end-relational theory if this difference could be explained as due to the interference of the ‘in order that...’ modifier.

¹⁵ See for example Thomson (2007), Harman (1973).

Compared against the ‘must’ of normative necessity and the ‘may’ of normative possibility, it is not difficult to locate the function of the normative ‘ought’. If I *must_e* φ (i.e., in order to bring it about that e) then it follows that I *ought_e* to φ , but *ought_e* doesn’t entail *must_e*. If I *ought_e* to φ then I *may_e* φ , but *may_e* doesn’t entail *ought_e*. ‘Ought’, therefore, is weaker than ‘must’ but stronger than ‘may’.¹⁶ Intuitively, I instrumentally ought to φ just in case it is *better*, with respect to bringing it about that e , that I φ than that I do any of the actions I could perform instead: ‘ought’ is essentially comparative. As we are seeking a reductive analysis of ‘ought’ into nonnormative terms, however, we need to reduce the concept of ‘better’. Fortunately in the instrumental case this is not difficult to do. An action φ is an instrumentally ‘better’ means to an end than another, ψ , just in case φ is more likely to lead to that end than is ψ —or so I shall argue.

It is arguable that the ordinary ‘ought’ also concerns probability, providing a promising parallel. Arguably both ordinary and normative ‘oughts’ concern the greater likelihood of some state of affairs over the likelihood of the members of a comparison class \mathfrak{R} . I propose the following account of the ordinary ‘ought’:

Modal Ought: It ought to be the case that $p = (\text{df.})$ It is more likely, given circumstances C , that p than that any member of \mathfrak{R} obtains.

Hence (13) is true just in case it is more likely given C that it will rain tomorrow than any relevant alternative. (In this case it is most plausible, in a temperate climate, that the only relevant alternative in the comparison class is that it *doesn’t* rain tomorrow). I suggest that this is a good first approximation at an analysis of the ordinary modal ‘ought’, as quantifying probabilistically over alethically or epistemically possible worlds.¹⁷ However, contemporary work in linguistics makes it clear that defining ordinary ‘ought’ is a challenging task, and I expect that the correct account will differ in some significant ways from *Modal Ought*. My present purpose is merely to demonstrate that a decent first approximation vindicates the Third Thesis. This would give us reason for optimism that the correct analysis will also do so.

If we follow the treatment of ‘must’ and ‘may’ above, we will define ‘ought_e’ in terms of *conditional probability*:¹⁸

¹⁶ See also Sloman (1970), Sæbø (2001), von Fintel and Iatridou (2006a), (2006c), Copley (2006).

¹⁷ One possible fruit of the end-relational analysis of ‘ought’ is an elegant explanation of the distinction between the ‘objective’ and ‘subjective’ normative senses of ‘ought’ (i.e., the distinction between ‘You ought to act on your actual reasons’ and ‘You ought to act on what you believe to be your reasons’). This analysis prompts the thought that objective normative ‘ought’ is an alethic modal, while subjective normative ‘ought’ is an epistemic modal. I explore this suggestion further in Finlay (2007b).

¹⁸ Some issues arise here from Kolmogorov’s classic formulation of conditional probability as a ratio of unconditional probabilities. This has the consequence that the probability of p is undefined if the probability of C is nil, which has counter-intuitive implications. It seems, for example, that we can sensibly assert that certain things ought to be done, given some actually impossible circumstances. I think we should reject Kolmogorov’s formula for precisely this reason, as some philosophers working on probability recommend: it sometimes makes sense to talk about what is probable given some impossible circumstances.

Ought_e: In order that e it ought to be the case that $p =$ (df.) It is more likely, given circumstances C including its being the case that e , that p than that any other member of \mathfrak{R} obtains.

The class \mathfrak{R} here is the comparison class of relevant alternatives to p : i.e., the class of propositions (states of affairs, events, actions) such that they could obtain (occur, be performed) in place of p . Note that there are no normative terms in the definition (assuming that its notions of *relevance* and *likelihood* are not normative)¹⁹; if it is adequate, we have succeeded in reductively analyzing instrumental ‘ought’.

There is a serious problem, however; *Ought_e* seems to get things the wrong way around. We found that the instrumental ‘ought’ concerns the comparative probability of e given that p . My analysis addresses rather the comparative probability of p given that e . What someone ought instrumentally to do in some situation should not depend upon what they happen to be most likely to do in that very situation, but *Ought_e* seems to imply just this.

This problem can be solved in the following way: ‘in order that...’ functions partly to restrict C , blocking the troublesome implications. When we utter hypothetical imperatives, we are concerned with which action open to a person would be most likely to lead to the end, and are not interested in the relative probabilities of the person choosing that action.²⁰ As those probabilities are irrelevant to the modal facts of interest to us, we exclude them from the initial background circumstances C_0 . If we then supplement these circumstances, in which each alternative is equally likely, with the circumstance that e does obtain (yielding C , the union of C_0 and e)—which I have claimed to be the main function of ‘in order that...’—then the action that, if performed, is in C_0 most likely to bring it about that e , will be the action that is most likely in C to be the one performed. Given only the fact of success, the most effective means is the most likely.

Is *Modal Ought* approximately correct? It is often claimed that ‘oughts’ in sentences like (13) have a *predictive* meaning; they function *performatively*, expressing predictions. However, this view struggles to deal with acceptable sentences like

(15) It ought to have rained today, but it didn’t.

(16) It ought to rain today, but it won’t.²¹

My present suggestion is that semantically, ‘ought’ is always probabilistic. This probabilistic ‘ought’ is however sometimes *used* predictively; we can make

¹⁹ I may be challenged on both points. (1) By the ‘relevance’ of an alternative, however, I have in mind merely conversational relevance, which I take to be definable without appeal to further normative concepts. (2) On some *subjective* interpretations of probability, it turns out to be a normative concept. Indeed, if probability is interpreted in terms of what an agent ought to believe, then my analysis of ‘ought’ turns out to be circular. Although I do not offer an interpretation of probability in this paper, I must reject such subjective views, which in my judgement put the cart before the horse.

²⁰ One could argue for this as follows: instrumental ‘ought’ statements presuppose the point of view of giving (at least hypothetical) advice. When a person advises another, he pragmatically assumes that the other’s decision is entirely in the balance.

²¹ The envisaged scenario is one in which the speaker believes herself to have some special insight into the future: prescience, divine revelation, or the like.

predictions by asserting facts about probabilities—and sometimes *used* normatively; we can advise or evaluate behaviour by asserting facts about its prospects for success.

A word is needed about the detaching problem. Originally I suggested that the end-relational theory might escape the problem because ‘in order that...’ is not obviously the formula for the antecedent of a conditional. But it turns out that hypothetical imperatives are conditionals after all: they are *probability* conditionals.²² In effect, ‘in order that *e*, it ought to be the case that *p*’ is equivalent to ‘Given that *C*, and that it is to be the case that *e*, then it ought to be the case that *p*.’ But detaching will nonetheless not trouble this account, for an abundance of reasons. First, the antecedent doesn’t concern an agent’s motivations, so the consequent doesn’t detach from facts like those. Second, the detached ‘ought’ is not normative, but merely probabilistic. Third, these *probability* conditionals do not allow detaching anyway, because the probabilities in the consequent only hold relative to their conditions. Fourth, whether the antecedent obtains is not of relevance to our interest in these propositions, which rather addresses comparative probabilities given merely possible circumstances. By asserting propositions concerning probabilities that are conditional on counter-to-fact circumstances, we succeed in giving practical advice and evaluating actions. To even think about detaching the consequent is therefore to miss the point.²³

The instrumental normative ‘ought’, like its kindred instrumental normative terms, can plausibly be identified with its ordinary modal counterpart under an ‘in order that...’ modifier. This is a significant success for the end-relational theory when we consider that its rivals currently in favour among metaethicists all hold that the ordinary ‘ought’ is something distinct. The theory thus has a significant advantage over its rivals in being able to unify more of the data.

2 Complications

Some difficulties and objections must be addressed before we can consider the case for the end-relational account of the instrumental use of ‘ought’ satisfactory. First, there is another way of interpreting the function of ‘in order that...’ On this reading, we use ‘in order that *e*’ simply to explain or clarify the reason or payoff for the action we say ought to be performed. Suppose I assert (4), ‘Max ought to mingle with the crowd,’ and overhearing me, Clifford demands to know *why* Max ought to do that. I might then respond with

(17) Max ought to mingle with the crowd *in order to evade arrest*.

This sentence differs from (6) just in that the ‘in order to...’ qualifier has moved from the beginning to the end. While there is no heavy-duty significance in the difference, (6) gives a slight presumption to the end-relational reading, while (17) gives a slight presumption to the following rival reading.

²² Not all probability conditionals need involve conditional probabilities, but these do.

²³ For a more detailed examination of this solution to the detaching problem, see my (2007b).

Plausibly, in (17) the ‘in order that...’ functions simply to explain why it is the case that Max ought to mingle with the crowd, and does not place any conditions on or modify the meaning of the ‘ought’ claim. If this is right, then (17) claims merely that Max ought to mingle with the crowd *because* by doing so he can evade arrest, and is only true if (4): Max ought to mingle with the crowd. If it is not true that Max ought to mingle with the crowd, then it is not true that he ought to do so in order to evade arrest, or for any other end. On the basis of this reading we might opt for a theory on which this normative ‘ought’ is not qualified in meaning (i.e., it is unsubscripted), expressing a primitive normative operator, or perhaps meaning ‘the thing we have most reason to do’. If ‘in order that...’ sometimes operates in this way, it is arguable that it always does.

How are we to adjudicate between these theories? The primitivist interpretation might seem to have some theoretical advantages over the end-relational reading; on its simple semantics ‘ought’ is always just ‘ought’, and not ‘ought_{e1}’, ‘ought_{e2}’, etc. However, although the end-relational theory recognizes a multiplicity of ways ‘ought’ is relativized, it also gives a universal semantics for ‘ought’ itself. ‘Ought,’ on this view, is no more semantically ambiguous than attributives like ‘real’, comparatives like ‘big’, or indexicals like ‘here’, which despite their complex interaction with context are not difficult to interpret.

I grant that ‘in order that...’ is sometimes used in this other way, to explain the grounds for an ‘ought’ rather than to qualify or condition it.²⁴ The evidence is compelling. But there is also compelling evidence that on many occasions ‘in order that...’ functions rather in the way my theory describes. Suppose Clifford gets me to agree that (4) is false; Max ought not to evade arrest, and hence ought not to mingle with the crowd, since he has a moral obligation to face the music for his crimes. But I am interested in the technical question of how he might evade arrest, and so I press Clifford, asking, “But *in order* to evade arrest, ought he not to mingle with the crowd?” Intuitively there is a question here that has not yet been settled in our conversation, and Clifford may legitimately reply, “Yes, I suppose that in order to achieve that goal that is what he ought to do.” It is too constraining on ascriptions of instrumental oughts to require that an agent first ought to be pursuing the end.

A second objection challenges my claim that being instrumentally better for bringing it about that *e* is purely a matter of *reliability*, or the likelihood of *e*. Surely there are many different ways in which a means can be better, including speed, safety, cost, efficiency, enjoyability, and morality. The surest way for Max to evade arrest might be to shoot himself in the head, but we wouldn’t thereby conclude that this is what he *ought* to do in order to evade arrest.

There are different ways of evaluating means, only one of which is instrumental evaluation. To judge a means *morally* best, for example, is different from judging it instrumentally best. Much the same can be said of considerations of safety, cost, efficiency and enjoyability.²⁵ We have multiple ends and desires, and when we

²⁴ This is not to concede that ‘ought’ is sometimes not end-relational, but merely that I have yet to explain how the end-relational theory can accommodate these uses.

²⁵ Speed may seem problematic, but we can accommodate this by being more careful about including temporal information in the specification of an end. If your purpose is indifferent as to the time of realization, a speedier means has no greater instrumental value.

evaluate a means to one end, others often sit in judgement too. Such judgements will not be strictly instrumental evaluations, but rather decision-theoretic evaluations.²⁶ Greater care in specifying the exact content of the end may also resolve many of the problematic cases. Arguably, for example, by saying (6) I really mean to express

(18) In order that Max evades arrest *unharmedly*, it ought to be the case that he mingles with the crowd.

This seems right, given that when I utter (6) I presumably have in mind Max's escaping unscathed. The objection fails to show that instrumental value is not analyzable in terms of probability.

However, this genuinely is a *complication*. For what we say we ought to do *in order that e* is not sensitive only to what is instrumentally valuable for the obtaining of *e*. Doesn't this disprove *Ought_e*? I suggest two replies. First, we've acknowledged that 'in order that...' sometimes has an explanatory rather than a qualifying function. Noninstrumental considerations in such cases present no problems for *Ought_e*, because it isn't supposed to cover them. Second, we can manipulate the class of relevant alternatives \mathfrak{R} to get the right results. So far I have treated \mathfrak{R} as consisting of the actions the agent can perform instead. But plausibly what counts as a 'relevant' alternative can depend on context and interests, so that (e.g.) immoral or self-destructive alternatives might not be included in \mathfrak{R} . We can consider this a clarification rather than a revision of \mathfrak{R} if we adapt our conception of what might be meant by talk of what acts an agent 'can' perform.

A third objection suggests that the theory commits the 'naturalistic fallacy' by mistaking a statement of necessary and sufficient conditions for a statement of identity. Supposing that necessarily it is the case that p_1 (in order that *e* A ought to φ) if and only if p_2 (for any action ψ that A could perform instead of φ , given that *e* obtains it is more likely that A φ s than that A ψ s), it remains possible that p_1 and p_2 are different propositions and that the meaning of the instrumental 'ought' is not given by *Ought_e*. Necessary coextension is not sufficient for identity or analysis.

Once necessary coextension is granted, however, the onus is on the skeptic to motivate doubts that we have not successfully analyzed the concept. This is the job of Moore's 'open question' challenge. If the analysis is successful, then it is claimed that the following question should seem trivially answerable in the affirmative ('closed') by virtue of an understanding of the meaning of its component words alone:

OQ: Is it the case that all and only those actions that are more likely than any alternative to result in its being the case that *e* are those actions that one ought in order that *e* to do?

Passing over the fact that contrary to the objection, analytic truths *can* be nontrivial and contestable, I venture my own judgement that *OQ* is indeed a closed question. I cannot conceive of how some action could be something that I instrumentally ought

²⁶ The simplest way of interpreting these friendly to the end-relational theory is as judgements of which means ought to be adopted *in order that expected value is maximized*. The technical notion of expected value presents no problems for naturalistic reduction.

to do for the sake of some end, if doing it is not more likely to result in that end than something else I could do—and I cannot conceive of how some action could *not* be something that I instrumentally ought to do for the sake of some end if doing it is more likely to have that result than anything else I could do (other than by virtue of my other ends, as discussed above).

My opponent must argue here that there is something in the concept of ‘ought’ that is missing from *Ought_e* or *vice versa*. One might say that since *Ought_e* is reductively naturalistic, it lacks the crucial element of normativity. However this just begs the question against my reductive naturalistic project. It may be claimed that *Ought_e* lacks some of the distinguishing features of the normative ‘ought’. But what features are these? I have argued that *Ought_e* captures several of the distinctive features of the normative ‘ought’. A common suggestion is that ‘ought’ expresses some attitude like approval or endorsement. It is true that on *Ought_e* it does not intrinsically do this, but it is implausible that the instrumental ‘ought’ does either; I need not express my approval of poisoning to judge that in order to murder without detection one ought to poison one’s victim.²⁷ Another common suggestion is that ought-sentences function to direct or advise, rather than to describe facts. However to infer directly from how sentences are used to what they mean is to commit the speech act fallacy, and it is easy to see how and why probabilistic, end-relational ought-sentences would be used to direct or advise. Further candidates would have to be evaluated as they are suggested, but I do not see any danger that such a feature can be found. On the other hand, the theory’s ability to give a unified treatment of the instrumental and ordinary ‘oughts’, and the finding that modals in general function normatively when placed under an ‘in order that...’, give us strong reasons to think that there is no ‘naturalistic fallacy’ here.

3 The categorical ought

“This is all very well,” someone might say, “but it is beside the point. Instrumental ‘oughts’ were never the philosophically troubling case; it is rather the *categorical* or noninstrumental ‘ought’ that has proven resistant to reductive naturalistic analysis, and it is implausible that *Ought_e* gives the meaning of *that*.” It is indeed the ‘categorical’ character of certain ‘oughts’ that has led many to deny that those ought-sentences could report natural facts.²⁸ Wittgenstein expresses the view most memorably:

I can only describe my feeling by the metaphor, that, if a man could write a book on Ethics which really was a book on Ethics, this book would, with an explosion, destroy all the other books in the world. Our words used as we use them in science, are vessels capable only of conveying meaning and sense,

²⁷ Expressivists can reply that in the case of instrumental ‘oughts’, I would only be expressing my approval of poisoning as a means to committing a successful murder. But once we start qualifying attitudes like this, it is not clear that they are not simply to be identified with beliefs of some kind.

²⁸ Wittgenstein (1997), Mackie (1977), Joyce (2001).

natural meaning and sense. Ethics, if it is anything, is supernatural and our words will only express facts... I said that so far as facts and propositions are concerned there is only relative value and relative good, right, etc... The right road is the road which leads to an arbitrarily predetermined end and it is quite clear to us that there is no sense in talking about the right road apart from such a predetermined goal. (1997, p. 67)

Philosophers who reject metaethical naturalism are however divided about whether categorical normative language expresses *nonnatural* facts, tries but fails to express facts, or doesn't play the fact-stating game at all.

My Fourth Thesis is that *Ought_e* gives the meaning of 'ought' in *all* its normative uses:

Fourth Thesis: Every normative 'ought' is simply the ordinary modal 'ought' under an 'in order that...' modifier.

(Correlatively, *Modal Ought* gives the meaning of 'ought' in every use). This will seem unpromising at the outset, but I ask only that the end-relational theory be indulged its own attempt at drawing this particular sword from its stone. Given its unification of instrumental and nonnormative 'oughts', if the end-relational theory can accommodate the data here then we have ample justification for accepting *Ought_e* as the correct analysis of 'ought' in its categorical use. I shall recognize the existence of a categorical *use* of 'ought', but deny the existence of a distinct categorical 'ought'. We first need to clarify the nature of a categorical use of 'ought'. My strategy will then be to investigate what sense can be made of categorical uses on the assumption that the end-relational theory provides a correct semantic analysis. If the phenomena are just what we could expect on the assumption that a theory is true, then they cannot pose any objection to that theory.

A superficial feature of allegedly categorical 'ought' sentences is that when we utter them, 'ought' is not commonly preceded or succeeded by 'in order that...', or any similar expression. This first observation about categorical use generates a naïve objection to the Fourth Thesis: in categorical uses of 'ought' we don't *say* 'in order that...', so we can't mean to assert end-relational propositions. This is a bad objection, of course, because we often neglect to say 'in order that...' even when we are clearly using 'ought' instrumentally.²⁹ If I say to Clifford, "Suppose Max is trying to evade arrest. Then what ought he to do?" Clifford may respond, "He ought to mingle with the crowd." This response implicitly assumes the modifier, 'in order that he evades arrest.' Often the sentences we utter are elliptical for more complex sentences, because often propositional content is supplied by context and is therefore left implicit. This linguistic observation does not prove that there are any normative uses of 'ought' that are not correctly analyzed by *Ought_e*. I claim:

Fifth Thesis: Normative 'oughts' without explicit 'in order that...' modifiers always contextually presuppose such modifiers.

²⁹ See also von Fintel and Iatridou (2006a).

I shall argue that this contextual presupposition is in fact an essential feature of categorical use.

The evidence for a distinct categorical use of ‘ought’ comes from more than merely features of surface grammar, however; Clifford’s use of ‘ought’ would not be considered categorical. Categorical ought-claims are sometimes characterized as *unconditional*: as involving an ‘ought’ that is not qualified by any (explicit or implicit) conditions—or, at least, by any conditions involving what the relevant agent desires or intends. But consider (3), ‘If you want to smoke a cigarette, you ought to chew gum,’ and

- (19) If you have an urge to drown your child in the bathtub, you ought to seek psychological help.

These are categorical uses of ‘ought,’ although they involve conditions concerning what the agent desires or intends.³⁰ The intended characteristic of categorical use is better captured as involving an ‘ought’ that is not ‘instrumental’ in the following sense: the action’s having instrumental value for some end desired or intended by the agent is not a necessary or sufficient condition for the proposition’s truth.

Against my Fourth Thesis (that *Ought_e* accounts for all normative ‘oughts’) it might therefore be objected that the truth conditions for some ought-propositions are not sensitive to what the agents involved in those propositions desire or intend. This is true, but it poses no problem for the end-relational theory. In abandoning the ‘hypothetical imperative’ formulation for the ‘in order that...’ formulation, we abandoned any claim that the relevant ‘oughts’ were contingent upon an agent’s desires or intentions (for this reason it may be misleading to label end-relational propositions ‘instrumental’). It can be true that (6), in order that Max evades arrest he ought to mingle with the crowd, even if Max has no desire or intention to evade arrest. In order to constitute a genuine objection to the end-relational theory, categorical use would have to involve an ought-proposition that is noninstrumental in the sense of being independent of instrumental value for *any* possible end.

We have yet to locate the real thrust of the objection, however, because assertion of (6) in such circumstances would also not be considered a categorical use of ‘ought’. Consider the following paradigms of categorical ought-sentences:

- (20) Citizens ought not to condone their government’s practicing torture;
 (21) Grover ought to brush his teeth more often;
 (22) You ought not to believe a contradiction.

What makes these plausible candidates for categoricity? One characteristic beyond those already canvassed is their practical significance. A speaker who utters (20)–(22), unlike someone uttering (6), presumably *endorses, recommends, or prescribes* the agent’s compliance. Categorical uses of ‘ought’ convey a practical import which merely end-relational ‘ought’ claims as such lack.

³⁰ It might be objected that the proper form of these propositions is: O(if p, then q)—in which case there are no conditions on the ‘ought’ itself. However if we take this route we cannot differentiate ‘hypothetical’ from categorical ought-propositions, since the former may be represented as O(if A desires that e, then A φs). We have lost a contrast that we wanted to keep.

Might a typical use of (20)–(22) be end-relational? Of each it is true that, if asked for an explanation of its truth, we could furnish an end to which the prescribed behaviour contributes. This allows us to suggest ‘in order that...’ modifiers as follows:

- (20′) [In order that suffering not be inflicted on innocent people] Citizens ought not to condone their government’s practicing torture;
- (21′) [In order that he preserves his health] Grover ought to brush his teeth more often;
- (22′) [In order that you avoid irrationality] You ought not to believe a contradiction.

These sentences are all plausibly true; might they also be plausible interpretations of what someone might mean to assert on some particular occasion of uttering (20)–(22)? It is highly unlikely that *every* use of (20)–(22) is elliptical for (20′)–(22′), but there are indefinitely many alternative ends, providing plausible interpretations of these other utterances. I suggest that if (20)–(22) are not true in virtue of their relation to these specified or alternative ends, then it is not immediately obvious that and how they are true at all. On the hypothesis that (20)–(22) are elliptical for claims like those expressed by (20′)–(22′), there is a ready explanation available for their special practical significance. In each case, the end is something that we find important or that matters to us (assuming a basic altruism). If we care about the end, then the ‘ought’ will have a particular seriousness to us; it concerns what is most likely to promote something that matters to us. If this is what categorical use of ‘ought’ amounts to, then it is perfectly compatible with the end-relational theory. ‘In order that I don’t kill you, you must come with me’ may be end-relational, but this ‘must’ is not lightly ignored.³¹

This treatment of (20)–(22) invites two objections that I must momentarily set aside. First, it raises the question of *ultimate* ends. I accommodated each sentence within the end-relational theory by furnishing an end to which the prescribed behaviour is plausibly a means. Consider, therefore, assertion of

- (23) Suffering ought not to be inflicted on innocent people;
- (24) Grover ought to preserve his health;
- (25) You ought to avoid irrationality.

It appears harder to find a plausible ‘in order that...’ modifier for each of these, and even if we can, the same problem will just reappear for those further ends. It seems that the end-relational theory cannot account for cases like these. I will argue, however, that ultimate ends are merely a limiting case for the theory. Second, it may seem that any proposed interpretation of (20)–(22) along the lines of (20′)–(22′) will be too specific to accommodate deliberation and disagreement over their truth. When we address the question of whether or not a government ‘ought’ to practice torture, many different ends may bear on the answer. It may seem implausible, then, that deliberation and disagreement are focused narrowly on a single end. My

³¹ The end may variously be important to any of the speaker, the audience, and the agent, yielding different kinds of significance that an ‘ought’ claim may have.

solutions to both of these problems emerge from my treatment of a further account of the nature of categorical use.

The serious challenge from categorical use arises from considerations about *pragmatics*, involving the circumstances in which use of ought-sentences is conversationally appropriate and useful. Here we find apparent evidence that we use 'ought' to assert normative propositions that are not qualified by any end. This comes from two sources: (i) our ability to communicate normative claims without any explicit 'in order that...' modifier, and (ii) the practical significance that these normative claims seem to have. I have already argued, of course, that neither of these features *per se* present problems for the end-relational theory. But this was on the assumption of certain contexts, which in genuinely categorical use need not obtain. I examine these problems in turn.

We need not explicitly utter 'in order that...' to communicate an end-relational proposition, I observed, provided that the relevant ends are salient in the conversational context, such that our audience can recognize the intended modifier. But in their familiar categorical use (e.g., in moral discourse) we utter unqualified ought-sentences like (20)–(22) when no relevant ends have been introduced into the context. It is therefore unclear how we could successfully communicate an end-relational proposition in these circumstances, and therefore how such speech acts could be intelligible and hence conversationally appropriate if they are genuinely end-relational. The fact that these categorical uses of 'ought' are nonetheless intelligible is therefore evidence that they express propositions involving a categorical 'ought'. Call this the *communication* problem.

An end-relational 'ought' will have practical significance for an agent, I observed, provided that (and to the degree to which) he desires or intends the relevant ends. But in their familiar categorical use, a speaker behaves as if 'ought' claims have practical significance for the agent regardless of whether he desires or intends any plausible candidate for a relevant end; this is reflected in the fact that it is conversationally appropriate, and even deemed important, to address categorical uses of 'ought' to the agents themselves. In categorical uses, like typical utterance of (20)–(22), a speaker conveys some form of demand, prescription, or advice about how the agent is to act, which is not contingent on the agent's desiring or intending any end such as those proposed in (20')–(22'). For example, whereas I am likely to withdraw or refrain from addressing to Max the unqualified assertion, 'You ought to mingle with the crowd' once I accept that he has no thought of evading arrest, I have no such disposition to withdraw the unqualified assertion of (21), 'Grover ought to brush his teeth more often,' once I accept that Grover has no desire or intention to preserve his health.

In its categorical use, 'ought' seems to have an essential prescriptive or advising function. But it is difficult to see how this function could be performed by speech acts that merely report that some behaviour stands in an instrumental relation to some end, an end towards which the agent happens to be indifferent. An agent might fully appreciate and be indifferent to how his behaviour is detrimental to the end, but his interlocutors may nonetheless continue to insist to him that he *ought not* to behave that way, as if this is something he still fails to recognize and that could

potentially influence him. Call this the *significance* problem. Here different theories of the meaning of ‘ought’ seem to find vindication; that these ‘oughts’ have an essential function like commanding, directing, or advising an agent what to do, or that they report rational requirements. Gibbard, for example, suggests that ‘ought’ expresses what a speaker hypothetically plans to do in the circumstances in question (2003).

The communication problem and the significance problem together constitute the real challenge from categorical use to the end-relational theory. They arise from pragmatic features of how we use ought-sentences in conversation, and I shall now argue that there is a common solution for both problems, based on the conversational principles first explored by Paul Grice.³² For the sake of argument, suppose that the Fourth Thesis is correct; *Ought_e* gives the meaning of *every* normative ‘ought’. Our question is: what sense, if any, would an audience then be able to make of a categorical use?

As we observed, end-relational ‘ought’ sentences can appropriately be uttered without explicit ‘in order that...’ modifiers in case the relevant end is salient in the context and readily identifiable. When we encounter an unqualified ‘ought’, therefore, we look first to the context, which can furnish an end in a number of ways. (i) An end may have been explicitly introduced into the conversation, as in my discussion with Clifford. By hypothesis this is not true of the categorical case. (ii) In some cases, the subject of the ‘ought’ claim yields a salient end. For artifacts and other objects with a clear function (such as clocks, referees, bombs, and guide dogs) we can talk about how they ought to behave without explicitly qualifying the ‘ought’.³³ Some philosophers attempt to accommodate a categorical ‘ought’ as based on the *human* function. These accounts are problematic not only because they make implausible claims about human nature, but also because they struggle to accommodate the practical significance of the categorical use of ‘ought’; facts about human nature just don’t appear to have the importance for every person that moral facts are presented as having. (iii) Wherever ‘oughts’ are addressed to persons in the mode of advice, or with the purpose of directing behaviour, we might suppose that they assume the agent’s own ends. But while categorical uses often seem to share this mode and purpose, by definition they are not (in this sense) ‘instrumental’ uses of ‘ought’. (iv) The *default* assumption, when none of the previous conditions obtain, is that an end is salient in case it is, in that situation, the shared, urgent concern of all the conversation’s participants. Again, this is not true of the categorical case.

All these cues fail us when confronted with a categorical use of ‘ought’. The speaker would appear to have failed to make an intelligible contribution to the conversation by consciously *flouting* the conversational maxim directing us to avoid obscurity—yet clearly he intends to communicate something. What might he mean

³² See Finlay (2004, 2005, 2006) for my previous treatments of these proposals.

³³ Many philosophers following Plato and Aristotle have sought to generalize this to a functional theory of all ‘oughts’ or some other normative term (Foot (2001), Thomson (2007), Bloomfield (2001)). The end-relational theory, however, can accommodate these uses as well as others that functional theories struggle to explain, because it doesn’t share the commitment to the claim that ‘ought’ is based in a thing’s nature).

to communicate? By flouting a maxim, Grice observes, a speaker can *exploit* it (1989, p. 30). An audience is able to infer from an utterance that the speaker presupposes the conditions under which that utterance would be conversationally appropriate. In our case, therefore, an audience can infer that the speaker presupposes that the relevant end *is* salient in the context. Therefore, I suggest, the audience is able to infer the speaker's presupposition of the *default* circumstance in which the relevant end is salient in the context: that the end is an object of shared, urgent concern for all present.

However, in paradigms of categorical use this presupposition is false; there are no plausible candidates for a relevant end that matters both to speaker and audience. What is an audience to make of a speech act that presupposes a context that transparently does not obtain? Here we encounter a *rhetorical device*.³⁴ To speak in a way that supposes something to be true of your audience that is clearly controversial or false is a way of expressing a demand that it *be* true of them; consider 'In this family, we do not belch at the dinner table!' and 'You *will* come here!'³⁵ By categorical use of 'ought', therefore, a speaker expresses the demand that his audience share his concern for the relevant end, and consequently for the behaviour at issue.

On the supposition that the end-relational theory is correct, therefore, a speaker may through categorical use of 'ought' (*omitting* specification of a non-shared end) express, by virtue of these conversational principles, the demand that his audience have concern for that end and that they act accordingly. I propose:

Sixth Thesis: Categorical uses of 'ought' are rhetorical uses of the end-relational 'ought' based on ellipsis.

On this proposal, the characteristic features of categorical use are understood largely as the expressivists suggest, but as a pragmatic rather than the semantic function of 'ought'. I shall now suggest that categorical uses of 'ought' are constituted by not one, but a family of different rhetorical uses, falling into two categories. In one, the audience can identify the end despite its not being in (the forefront of) the context.³⁶ In the other, they are unable to identify any relevant end.

An audience may be able to identify an appropriate end despite its being missing from (the foreground of) the shared context in at least the following three ways. First, the end might be discernable from the speaker's own known concerns. Second, a categorical use of 'ought' can invoke a social institution based on this rhetorical device, whereby there are certain ends that are socially 'expected' of agents: a *morality*. Where such an institution exists, an audience is able to glean from content and categorical use of an ought-sentence that it assumes qualification

³⁴ See my (2004), (2005), (2006) and also Barker (2000).

³⁵ The very ambiguity of 'expects' appears to illustrate this phenomenon. To 'expect' (in the normative sense) that someone will do their duty is just to adopt a rhetorical stance by pretending to 'expect' (in the nonnormative sense) that he will. The rhetorical force seems to arise from challenging others to dare to contradict you in speech or behaviour.

³⁶ Once the rhetorical use is recognized, an audience has new places to look for a relevant end. Whether the right thing to say is that the relevant end is not in the context, or merely that it is less salient in the context, depends on fine issues about what conversational context is, which I will not address here.

by these moral ends. In my view this is often what happens when we make moral claims.³⁷

My third suggestion offers a preliminary answer to the problem of ultimate ends, previously set aside. Ought-propositions concerning ultimate ends may be merely a limiting case for the end-relational theory, because we can generate trivially true end-relational interpretations of (23)–(25) as the following tautologies:

- (23') [In order that suffering not be inflicted on innocent people,] Suffering ought not to be inflicted on innocent people;
 (24') [In order that he preserves his health,] Grover ought to preserve his health;
 (25') [In order that you avoid irrationality,] You ought to avoid irrationality.

These interpretations may seem quite implausible. While (23')–(25') are trivial, (23)–(25) seem to be potentially relevant, informative, and even important for communicative purposes. However, if the real conversational function of uttering (23)–(25) is to demand motivation towards the relevant ends (or at least the prescribed behaviour) rather than to convey their semantic content, then their significance is quite compatible with their being tautologous. We often find that communicative purposes can be served by asserting tautologies: consider 'A fact is a fact,' and 'It ain't over till it's over'.³⁸ With regard to the communication problem, tautologies may present an advantage. If the relevant ends aren't salient in any other way, there may be a default assumption that the intended end is identical with the 'means'. This interpretation of (23)–(25) is admittedly speculative, but the end-relational theory has other resources for accommodating ultimate ends, to which I now turn.³⁹

The second category may seem more problematic. How can a speaker communicate a demand for concern for some end, if the audience is unable to identify the relevant end? There are numerous ways of extending the account here. (i) A speaker may intend a particular end but fail to communicate which. What is communicated instead is that there is some particular end such that the speaker expects (i.e., requires) his audience to care about that end, and for the sake of which one ought to φ . Although the end is not identified, the speaker successfully expresses a demand that others be motivated to φ . (ii) The speaker may express an existential proposition: there are some important ends such that in order that those ends are realized it ought to be the case that p . A parent might mean this, for

³⁷ See Finlay (2007a) for a more careful presentation of this suggestion. It is important, I believe, to distinguish evaluative judgements with a peculiarly moral content from evaluative judgements made from a peculiarly moral stance. Here I mean the former, which accommodates the 'amoralist' but not the moral revolutionary.

³⁸ See Grice (1989, p. 33). Other examples include 'A lie is a lie,' 'It is what it is,' and 'I believe what I believe.'

³⁹ Unlike my other examples of informative tautologies, in these cases it isn't obvious that they are tautologies. This difference is adequately explained, I think, by the fact that they are not *explicitly* tautologous. But it is arguable that we do implicitly recognize their (implicit) tautologous character: there does seem to be something lame about taking a stand on such normative claims in conversation (e.g., 'You ought not to hurt people. It's just *wrong*.') We tend to feel that there is no point continuing to argue with a person who would deny this. I suggest this is because we understand implicitly that here there is no disagreement about facts, but rather a basic difference in attitudes. Under normal conditions, moral discourse presupposes some level of shared concern.

example, when imperiously telling his child, ‘You ought to eat your veggies!’ This option admits a form of *error theory*; a speaker might falsely believe or hope that there are such ends, without having any idea what they might be. (iii) A form of *fictionalism* might be true of some categorical uses of ‘ought’; the speaker *pretends* that there are some important and shared ends. Finally (iv) if the semantic and pragmatic accounts I have given of the uses of ‘ought’ are correct, then a parasitic use should be possible. Since end-relational ‘oughts’ are characteristically used to advise, demand, endorse, and criticize, an ought-sentence could be given a purely expressivist use to perform such speech acts, without regard to its semantics and without any intention to express propositions. (Consider, for example, how we use ‘How are you?’ to perform an act of greeting, or how an atheist might say ‘Thank God for that!’ to express relief, without regard for semantic content.) In this case, ‘ought’ is being used as a ‘word of mere mesmeric force.’⁴⁰

On the hypothesis that the Fourth Thesis is correct—the meaning of ‘ought’ in its categorical uses is given by *Ought_e*—it therefore turns out to be plausible that categorical uses of ‘ought’ express demands and attitudes just as expressivists claim. Because of this we can engage in categorical normative ‘disagreement’ with each other while uttering sentences that semantically express noncontradictory propositions, thereby expressing conflicting attitudes and demands. On its own terms, therefore, the end-relational theory can accommodate and explain the categorical use of ‘ought’. Given the ability to account for predictive, instrumental, and categorical uses of ‘ought’ with a single reductive semantic analysis, the end-relational theory has a strong case for being preferred to any account postulating a distinct, categorical ‘ought’.⁴¹

We can promote this account of the categorical use of ‘ought’ further by considering the following virtues. First, it accommodates and explains in a simple and plausible way the intuition that in its categorical use, ‘ought’ is ambiguous between different senses, having distinct uses pertaining at least to morality, prudence, epistemology, etiquette, and law. Consider

- (26) [In order that innocent people not be harmed] Citizens ought not to permit their government to practice torture;
- (27) [In order that he enjoys a life of maximal wellbeing] Grover ought to brush his teeth more often;
- (28) [In order that your belief set be as comprehensive and correct as possible] You ought to believe what you have good evidence for.

⁴⁰ Anscombe (1958).

⁴¹ There is a difference, however, between showing that the end-relational ‘ought’ *could* be used categorically, and showing that our actual categorical uses of ‘ought’ are end-relational. The latter has not been accomplished here, and a serious objection that I cannot attend to sufficiently is that since people don’t ordinarily take themselves to be asserting end-relational propositions when they utter ought-sentences, it is most unlikely that they are. I would insist on distinguishing sharply between what we mean by our words and what we think we mean, and defend the end-relational theory on this basis, but this is an issue that must wait for another time. This lack of opacity might seem especially damaging to my pragmatic account, which depends upon an audience having a firm grasp of the semantic rules of ‘ought’. I would argue that here also our competence turns on our intuitive grasp of—rather than any explicit theoretical understanding of—the meaning of our words.

- (29) [In order that one behaves in a manner that will not cause offense] One ought not to talk with one's mouth full.
- (30) [In order to comply with the law] One ought not to drive on the right-hand side of the road.

On the end-relational theory, there will be distinct 'oughts' of morality, prudence, epistemology, etiquette, and law simply in virtue of the fact that there are distinct moral, prudential, epistemological, social, and legal ends. Unlike many of its rivals, this theory is able to explain very simply how it can be the case that what we morally ought to do differs from what we legally or prudentially ought to do, for example.⁴²

Second, it provides an intuitively plausible explanation of *why* we ought (in whichever sense) to do the specific things that we ought to do rather than other things—an explanation that is behind the enduring appeal of teleological approaches to moral theory. Utilitarianism, for example, is animated by the perception that what we morally ought to do can be identified with what is most promotive of maximal human wellbeing impartially considered, and Kantianism, with what we must do in order that we are able to will our maxims as universal law.⁴³

Third, it possesses the chief virtues but lacks the chief vices of rival theories. Expressivists urge us to accept their theories on the ground that a descriptivist (naturalist or nonnaturalist) analysis of 'ought' cannot accommodate the essential practical significance of categorical normative judgement. If 'ought' merely signifies some kind of relation, for example, how could it possess its essential demanding character? But the end-relational theory is immune to this objection, because it is able to explain the demand as a rhetorical product of the categorical use of the end-relational 'ought'. Expressivism, on the other hand, struggles with the appearance, accommodated here, that ought-sentences express real propositions, representing real objective facts. Furthermore, expressivists are hard-pressed to explain *why* normative utterances would have the influence on our attitudes and behaviour that they do, while the end-relational theory offers a ready explanation.

4 Ought simpliciter

There is one remaining challenge for the end-relational theory, which maintains that all normative ought-sentences are relativized to particular ends (subscripted). What about 'ought' *simpliciter*: plain normative 'ought'? The theory denies that there is any such thing, which strikes many as implausible. In this section I endeavour to sweeten this pill.

⁴² By contrast, Thomson's univocal, functionalist semantics for 'ought' (2007) leads her to claim that when we say of a game of chess that a person 'ought' to move his pawn, then we mean that this is something that he *morally* ought to do. While this claim is not as absurd as it may at first appear, it is counter-intuitive. The end-relational theory will rather qualify this 'ought' with the modifier, 'in order to win at chess'.

⁴³ Of course, these monistic theories each struggle to accommodate all the data. But this may be merely because it is a mistake to define morality by just one end.

One reason for supposing that there is an ‘ought’ that transcends our ends is the commonplace talk of what we ought to do, ‘all things considered’. An obvious reading of this is as appealing to an ‘ought’ that remains once we have abstracted ourselves from all particular ends. But there are other, natural interpretations that are friendlier to the end-relational theory. If normative ‘oughts’ are indeed conditional probabilities, then it is natural to read the contrast between *prima facie* and all-things-considered ‘oughts’ as the contrast between probabilities conditional on partially specified circumstances and those conditional on fully specified circumstances. The ‘things’ to be considered therefore need not be ends. Alternatively, the quantifier may be restricted (for example, to the set of the agent’s ends, or the set of important ends). And even if the ‘things’ are ends, the ‘ought’ itself may yet be relative to a single end. It makes perfect sense to say, ‘Taking all my ends into consideration, I ought (in order that e) to φ ’.

Is there other evidence that we employ a concept of ought-simpliciter? We have seen that the fact that we utter ought-sentences without explicit qualifiers counts for nothing in light of the role of context and ellipsis. Neither can we take the utterance of ought-sentences when there is no relevant modifier in the context as evidence that they are not end-relational, given the rhetorical utility of such utterances on the assumption that all normative ‘oughts’ are end-relational.⁴⁴ The real basis for such convictions seems to be the perceived need for and possibility of (a) deliberation between our ends, and (b) critical reflection on them. Aristotle’s claim that deliberation is always of means is widely thought implausible. Finding that we have incompatible ends, we reflect on our options, weigh the ends, and decide what to do. Our conclusion, it is argued, must be a conclusion about what we *ought* to do; an ‘ought’ that must transcend those ends. Our ability to weigh our ends, it is further argued, is evidence that there is a common evaluative scale, transcending those ends, on which each can be weighed, such that it is possible to find one better than the other, and hence what we *ought* to pursue.

Is this decisive against the end-relational theory? It is possible to weigh an end on the scales of some *further* end, and so it can be suggested that weighing our ends involves turning to some further end to adjudicate between them. Some suggest that deliberation as such may have constitutive ends, such as acting in a self-intelligible way (David Velleman (2000)), or being true to the conception of your identity under which you value yourself (Christine Korsgaard (1996)). I find it more plausible that individual instances of deliberation are motivated by different ends. If Max hesitates before fleeing arrest to weigh the end of prolonging his freedom against the end of conforming with his moral obligations, he may do so with the aim of choosing the option that will result in the most happiness for him in the long term, or of choosing

⁴⁴ Neither is it a good objection that it is impossible to pin down a particular end to account for various utterances of ought-sentences made within a conversation. Context can change rapidly, from sentence to sentence or even clause to clause. Consider the following conversation, on a visit to a 3th grade classroom. A: ‘Carlos is very tall’. B: ‘How about his father?’ A: ‘No, his father is short.’ B: ‘So Carlos is taller than his father, then?’ A: ‘No, of course not—his father is taller.’ This conversation is quite coherent, but how can it be coherent that a short person is taller than a tall person? The answer, of course, is that A’s use of ‘tall’ with regard to Carlos is ‘tall for a 3rd grader’, while his use of ‘short’ for the father is ‘short for an adult’. The change in context is facilitated in this case by the change of subject and comparison class. Similar dynamics may account for rapid changes in the relativization of ‘ought’.

the action that he will regret the least, or of minimizing suffering for his loved ones. It is unclear how a person is motivated to deliberate at all if not by some desire or intention: the end-relational theory will identify the operative end with whichever end motivates deliberation.

It may be objected that such a regress of ends must terminate somewhere, whereupon the problem re-emerges when we have to weigh our ultimate and dearest ends. However this assumes that we only decide between ends on the basis of some dearer end, and this need not be the case: the structure of deliberation could be coherentist rather than foundationalist. If we do not simply prefer one ultimate end to another, then perhaps we choose between them on the basis of less dear ends.

Consider, however, the plight of someone who must choose between his only two ends, sacrificing one to the other (Cepheus, perhaps, choosing between saving Joppa, his kingdom, and saving Andromeda, his daughter). Mightn't he coherently address his dilemma by asking himself, 'Which ought I to choose?' The intended 'ought' cannot be merely relative to one end or the other, and neither can it be relative to any further end. The end-relational theory is committed to denying that there can be an intelligible normative question here.⁴⁵ While I am happy to bite this bullet, denying the existence of any normativity that transcends all ends, others may see this as an unacceptable cost. However, there may still be an intelligible *deliberative* question here. It seems false that deliberation is always aimed at determining what one *ought* to do. At least some deliberation seems aimed rather at determining what one *shall* do, which seems to be a separate matter. If deliberation was exhausted by determining what we ought to do, we would have a serious problem when, as for Buridan's ass, the weights appear even and so there is no right choice. To resolve such dilemmas it seems we *need* to have a capacity for deliberation about what to do, distinct from deliberation about what we ought to do. The champion of end-transcendent normativity must then explain why it can't be *this* capacity that we exercise when torn between ends. Where we deliberate over conflicting and incommensurable ends it may then be that we are not trying to determine what we ought to do, but merely to decide what to do.⁴⁶

Of course, it seems natural to use 'ought' when expressing these deliberative questions. But in the previous section we saw that the end-relational theory provides resources to explain how and why we might *use* 'ought' as if it involved an absolute normative relation. In this derivative, expressivist use, 'ought' will function as if it were an 'ought'-simpliciter, signifying the content of decision, or (in Gibbard's expression) the 'thing to do'. Proving that there is no 'ought'-simpliciter is too ambitious a goal for this paper, but I hope to have succeeded in showing that the end-relational theory has a coherent and defensible way of explaining away the evidence that there is such an 'ought'.

I conclude that in order to advance our comprehension of 'ought', we ought to consider the end-relational theory a serious contender, given (i) that it attributes to

⁴⁵ It is worth remarking that a psychology this sparse is barely imaginable as human, so we may reasonably be skeptical about the intuition that this person could genuinely ask himself this normative question.

⁴⁶ Gibbard (2003) explicitly denies any difference.

‘ought’ a single semantic contribution in every use, normative and nonnormative; (ii) that it is still able to give a simple explanation for why ‘ought’ seems to have different senses that can conflict without being contradictory; (iii) that it tracks a parallel function for other modals in the same contexts; (iv) that it is ‘reductively naturalistic’, offending no scruples against metaphysical, epistemological, or motivational ‘queerness’; (v) that it enables us to explain, by appeal to plausible general conversational principles, why normative language can have all the uses to which expressivists point, but (vi) it does so without incurring the semantic costs of expressivism. This is, I hope, some redemption for the much-maligned metaethical program of reductive naturalistic analysis.⁴⁷

References

- Anscombe, G. E. M. (1958). Modern moral philosophy. *Philosophy*, 33, 1–19.
- Barker, S. (2000). Is value content a component of conventional implicature? *Analysis*, 60, 268–279.
- Bloomfield, P. (2001). *Moral reality*. New York: Oxford University Press.
- Broome, J. (1999). Normative requirements. *Ratio*, 12, 398–419.
- Copley, B. (2006). What should ‘should’ mean? (manuscript).
- Dreier, J. (2006). Practical conditionals (manuscript).
- Finlay, S. (2004). The conversational practicality of value judgement. *Journal of Ethics*, 8, 205–224.
- Finlay, S. (2005). Value and implicature. *Philosophers’ Imprint*, 5(4), 1–20.
- Finlay, S. (2006). The reasons that matter. *Australasian Journal of Philosophy*, 84, 1–20.
- Finlay, S. (2007a). Too much morality. In P. Bloomfield (Ed.), *Morality & self-interest*. New York: Oxford University Press.
- Finlay, S. (2007b). What ‘ought’ probably means, and why you can’t detach it (manuscript).
- Finlay, S. (forthcoming). Against all reason? Skepticism about the instrumental norm. In C. Pigden (Ed.), *Hume, motivation and virtue*. MacMillan.
- Foot, P. (2001). *Natural goodness*. Oxford: Oxford University Press.
- Geach, P. (1991). Whatever happened to deontic logic? In P. Geach (Ed.), *Logic and ethics*. Dordrecht: Kluwer.
- Gibbard, A. (2003). *Thinking how to live*. Cambridge, MA: Harvard University Press.
- Grice, P. (1989). Logic and conversation. In P. Grice (Ed.), *Studies in the way of words*. Cambridge, Mass.: Harvard University Press.
- Hampton, J. (1998). *The authority of reason*. Cambridge: Cambridge University Press.
- Harman, G. (1973). Review of R. Wertheimer’s the significance of sense. *Philosophical Review*, 82, 235–239.
- Hill Jr. T. (1973). The hypothetical imperative. *Philosophical Review*, 82, 429–450.
- Joyce, R. (2001). *The myth of morality*. Cambridge: Cambridge University Press.
- Korsgaard, C. (1996). *The sources of normativity*. Cambridge: Cambridge University Press.
- Korsgaard, C. (1997). The normativity of instrumental reason. In G. Cullity & B. Gaut (Eds.), *Ethics and practical reason*. Oxford: Oxford University Press.
- Kratzer, A. (1977). What ‘must’ and ‘can’ must and can mean. *Linguistics and Philosophy*, 1, 337–355.
- Mackie, J. L. (1977). *Ethics: Inventing right and wrong*. Harmondsworth: Penguin.
- Sæbø K. J. (2001). Necessary conditions in a natural language. In C. Féry & W. Sternefeld (Eds.), *Audiat Vox Sapientiae: A Festschrift for Arnim von Stechow*. Akademie Verlag.
- Schroeder, M. (2004). The scope of instrumental reason. *Philosophical Perspectives*, 18, 337–364.
- Schroeder, M. (2005). The hypothetical imperative. *Australasian Journal of Philosophy*, 83, 357–72.

⁴⁷ I thank Kent Bach, Gunnar Björnsson, Matthew Chrisman, Bridget Copley, Matt Kaptein, Jeff King, David Manley, Jake Ross, Scott Soames, Gideon Yaffe, the Department of Philosophy at NYU, participants in my Fall 2007 graduate seminar, the PEA Soup bloggers, the *Phil Studies* referee, and especially Mark Schroeder, for their invaluable help in making this paper far better than it would otherwise be.

- Schroeder, M. (2007). Do oughts take propositions? (manuscript).
- Setiya, K. (2007). Cognitivism about instrumental reason. *Ethics*, 117, 649–673.
- Sloman, A. (1970). Ought and better. *Mind*, 79, 385–394.
- Thomson, J. J. (2007). Normativity. In R. Shafer-Landau (Ed.), *Oxford studies in metaethics*, Vol. 2. New York: Oxford University Press.
- Velleman, J. D. (2000). *The possibility of practical reason*. Oxford: Oxford University Press.
- von Fintel, K., & Iatridou, S. (2006a). What to do if you want to go to Harlem: Anankastic conditionals and related matters (manuscript).
- von Fintel, K., & Iatridou, S. (2006b). On the interaction of modals, quantifiers, and if-clauses (manuscript).
- von Fintel, K., & Iatridou, S. (2006c). How to say ought in foreign: The composition of weak necessity modals (manuscript).
- Wallace, R. J. (2001). Normativity, commitment, and instrumental reason. *Philosophers' Imprint*, 1(3), 1–26.
- Wedgwood, R. (2006). The meaning of 'Ought'. In R. Shafer-Landau (Ed.), *Oxford studies in metaethics*, Vol. 1. New York: Oxford University Press.
- Williams, B. (1981). Ought and obligation. In: B. Williams (Ed.), *Moral luck*. Oxford: Oxford University Press.
- Wittgenstein, L. (1997). Lecture on ethics. In: S. Darwall, A. Gibbard, & P. Railton (Eds.), *Moral discourse and practice*. Oxford: Oxford University Press.