

What *ought* probably means, and why you can't detach it

Stephen Finlay

Received: 29 November 2007 / Accepted: 10 July 2009 / Published online: 2 August 2009
© Springer Science+Business Media B.V. 2009, corrected publication 2022

Abstract Some intuitive normative principles raise vexing ‘detaching problems’ by their failure to license modus ponens. I examine three such principles (a self-reliance principle and two different instrumental principles) and recent strategies employed to resolve their detaching problems. I show that solving these problems necessitates postulating an indefinitely large number of senses for ‘ought’. The semantics for ‘ought’ that is standard in linguistics offers a unifying strategy for solving these problems, but I argue that an alternative approach combining an *end-relational* theory of normativity with a *comparative probabilistic* semantics for ‘ought’ provides a more satisfactory solution.

Keywords Semantics of ought · Detaching problem · Normativity · Deontic modals · Instrumental principle

1 Introduction

Certain intuitive normative principles raise puzzles about the meaning of ‘ought’ that continue to vex ethical philosophy. These principles take the form of conditionals, but detaching their consequents by modus ponens yields unacceptable results. Section 2 examines three principles and the *detaching problems* they present. I show that solving these problems forces us to postulate indefinitely many different senses for ‘ought’, which presents a further problem. How can the one word have so many different meanings? How can we manage to distinguish all these different senses, so as to be able to identify the relevant one in any case?

S. Finlay (✉)
University of Southern California, Los Angeles, USA
e-mail: finlay@usc.edu

This paper is partly remedial; I point out that the standard semantics for ‘ought’ widely accepted by linguists (following Kratzer 1977, 1981), of which ethical philosophy is still largely ignorant, provides resources that allow us to (i) accommodate and differentiate these different senses of ‘ought’ in a single semantic theory, and (ii) explain why the troublesome principles do not license detachment of their consequents. But I also demonstrate that a variant on the standard semantics (combining a *comparative probabilistic* [CP] theory of the semantic core of ‘ought’ with an *end-relational* [ER] theory of normative or ‘deontic’ content) can satisfy these and additional desiderata, providing an account that is also (iii) semantically simpler, (iv) more unifying, (v) better able to accommodate and explain the principles examined in Sect. 2, and (vi) metaethically *reductive*, analyzing ‘ought’ sentences without using normative terms. This theory seeks to advance our understanding of normativity itself, undermine skepticism about its existence, and cast light on why certain normative principles are true and not others.

Section 3 examines the standard semantics, Sect. 4 introduces the CP/ER approach, and Sect. 5 applies it to the principles from Sect. 2. My aim here is not to establish the truth of this CP/ER theory,¹ but merely to show that there is a theory of ‘ought’ that can satisfy all of (i)–(vi).

2 The detaching problem and the proliferation of senses of ‘ought’

It is common to distinguish many different senses of normative ‘ought’, including moral, epistemic, prudential, instrumental, rational, and legal senses, and even senses of chess, etiquette, cricket, etc. Many resist this fragmentation of the normative; for example Judith Jarvis Thomson gamely argues, in defense of the view that all practical ‘ought’s are univocally *moral*, that it is only true that some chess player ought to move his rook if he would otherwise be morally defective (2007). A strong case for ambiguity emerges from a ‘detaching’ problem that arises for many normative principles. This section explores the problem in some of its commonly discussed guises, and how the search for solutions leads to troubling proliferation in the senses of ‘ought’.

Consider the principle that perplexed Ewing (1953, p. 144f);²

Self-Reliance Principle: If an agent *S* believes that she ought to do *A*, then she ought to do *A*.

There is something intuitively right about this principle, which captures the platitude that one ought to follow one’s conscience.³ If Jorja believes that she ought to skip school and yet fails to (intend to)⁴ do so, then she or her behaviour is in some way

¹ I offer arguments for this in Finlay (2009, forthcoming a, forthcoming b).

² Also Dancy (1977); Gensler (1985); Broome (1999); Schroeder (2004, 2009); Kolodny (2005, p.512); Piller (2007).

³ Invoking conscience suggests a narrowly moral interpretation of the principle, however, which as we shall see has more general application.

⁴ The principle can be formulated either as requiring action or as requiring intention. Criticism of irrationality is better suited to the latter, since some ways of failing to act do not implicate one’s rationality. I shall

defective (*irrational*, some say).⁵ But there also seems to be something wrong with the principle. It appears to imply something implausible: that a person can never be mistaken about what she ought to do. By modus ponens, if a conditional and its antecedent are both true, we can *detach* the consequent. From the Self-Reliance Principle and Jorja's believing she ought to skip school, we should therefore conclude simply that Jorja ought to skip school. But surely a person can believe *falsely* that she ought to do *A*. If her belief is false, then it is not the case that she ought to do *A*. Hence the circumstance that Jorja believes falsely that she ought to skip school yields a contradiction; it is and it is not the case that she ought to skip school. This is the detaching problem for the Self-Reliance Principle; a solution must explain how it can be true that if someone believes she ought to do *A* then she ought to do *A*, without incurring these unpalatable consequences.

Consider next the

Subjective Instrumental Principle: If an agent *S* intends an end *E* and believes that doing *M* is the necessary means to achieving *E*, then *S* ought to do *M*.

This is a rather narrow instrumental principle—often the advisable means to some end is merely the *best* of several—but this focus is common in discussions of instrumental reason.⁶ The detaching problem here parallels the previous case. There is something intuitively right about this principle. If Jorja intends to skip school and believes that feigning illness is the necessary means, but fails to (intend to) feign illness, then she or her behaviour is in some way defective ('irrational'). But again, the principle appears to imply something implausible: that a person always acts as she ought when she successfully pursues her ends. By modus ponens, we could conclude simply that Jorja ought to feign illness. But surely it is possible that Jorja *ought not* to skip school, and so ought not to feign illness. The circumstance that Jorja intends to skip school, as she ought not, and believes of some action which she ought not to perform that it is a necessary means to skipping school, yields a contradiction; it is and it is not the case that she ought to feign illness.

In general, the detaching problem consists in a form of bootstrapping;⁷ these principles seem implausibly to make our attitudes self-legitimizing and therefore normatively infallible. Our normative beliefs make themselves true, and our intentions make their own executions right.

One obvious strategy for escaping the problem is to postulate that distinct senses of 'ought' are in play, so the contradictions are only apparent. But many philosophers

Footnote 4 continued

mostly ignore the difference, since the formulations are tightly linked. If the principle requires *S* to do *A*, then it plausibly follows that she is required to intend to do *A* as the means to acting. If it requires *S* to *intend* to do *A*, this is presumably because intending is a means to acting (i.e. this case is not akin to toxin puzzle cases, where the point of intending is separate from the point of acting).

⁵ Such a failure is not necessarily irrational. Many think, as I do, that there is nothing necessarily irrational about failing to try to act as you believe you *morally* (or prudentially) ought, for example.

⁶ Example Hill (1973); Gensler (1985); Broome (1999); Broome (2001); Broome (2002); Wallace (2001); Schroeder (2004, 2009); Setiya (2007). As before, I am ignoring a distinction between the requirement of intention and the requirement of action; see footnote 4.

⁷ Bratman (1987, pp. 23–27); Broome (2002, p. 96); Kolodny (2005, p. 512).

resist multiplying the senses of ‘ought’. Among the champions of a single practical sense of ‘ought’ (often interpreted as ‘has most practical reason’),⁸ we find two strategies. The first rejects the formulations of the principles above for formulations with *normativized* antecedents. The genuine Instrumental Principle, for example, prescribes only that if an agent *ought* to be intending an end then she ought to take what she believes to be the necessary means.⁹ The consequent ‘ought’s do detach when the antecedents are satisfied, but this only occurs when no contradictory ‘ought’ is true. However, this strategy fails to accommodate the ‘rational’ defects in failing to (try to) act as one wrongly believes one ought, or failing to take means believed necessary for one’s wrongly intended ends.¹⁰ Such behaviours seem to manifest precisely the defects that the Self-Reliance Principle and Subjective Instrumental Principle proscribe.

The second strategy accepts that these behaviours violate requirements of rationality, which it explains by appeal to *wide-scope* ‘ought’s.¹¹ On this view, these principles are properly understood as involving an ‘ought’ with scope over the conditional rather than the consequent; i.e. $O(A \rightarrow B)$ rather than $A \rightarrow OB$. The Self-Reliance Principle tells us that (e.g.) Jorja ought to make it the case that: if she believes she ought to skip school then she skips school. This is a requirement to make a disjunction true; i.e.

Ought (S does A or S doesn’t believe that she ought to do A).

We escape the detaching problem and bootstrapping because the normative claim that S ought to do A (e.g. Jorja ought to skip school) is not a consequent of this principle and so does not detach from it. It is consistent with this principle that Jorja ought *not* to skip school, because it allows that the only acceptable way for her to comply might be by ceasing to believe that she ought to skip school.

Similarly, the Subjective Instrumental Principle would tell us (e.g.) that Jorja ought to make it the case that (if she intends to skip school and believes feigning illness is the necessary means, then she feigns illness). It is the principle

Ought (S does A or (S doesn’t both intend the end E and believe doing A is the necessary means to E)).

We escape the detaching problem, again, because the claim that S ought to do A does not detach from this principle. It is consistent with this principle that Jorja ought not to feign illness, because the principle allows that the only acceptable way of complying might be to cease intending to skip school.¹²

Persuasive objections have been raised against wide-scope readings of these two principles, however.¹³ Wide-scope principles simply require that the conditional be

⁸ For example Broome (1999); Piller (2007); Thomson (2007).

⁹ Korsgaard (1997); Hampton (1998, p. 163); see Piller (2007, p. 63), for the case of the Self-Reliance Principle.

¹⁰ Elsewhere (forthcoming a) I argue that such instrumental irrationality isn’t even possible. But I don’t deny that if it were, it would be a defect in the will.

¹¹ Hill (1973); Dancy (1977); Broome (1999, 2001, 2002); Wallace (2001); Way (forthcoming).

¹² Or cease believing that feigning illness is the necessary means! (Schroeder 2004).

¹³ Schroeder (2004); Kolodny (2005); Setiya (2007); Piller (2007); see Way (forthcoming) for a defense.

true, and are indifferent to how this is accomplished. Faced with a conflict between one's actions and one's conscience, the wide-scope Self-Reliance Principle smiles as benignly upon changing one's conscience as upon changing one's behaviour. Faced with a conflict between one's attitudes towards the means and one's intention for the end, the wide-scope Subjective Instrumental Principle is indifferent to whether one gives up intending the end, comes to intend the means, or abandons the belief in the necessity of the means.

It can be observed in defense of the wide-scoping strategy that some extraneous normative principles may explain the impermissibility of satisfying the principles in the wrong ways, e.g. a proscription against revising beliefs for non-epistemic reasons. But this doesn't help, because we can still force a detaching problem for any wide-scope 'ought' by constructing scenarios in which satisfying the principle in the other way is blocked. One such scenario involves the agent's being justified in rejecting the alternative;¹⁴ if Jorja ought not to revise her false belief that she ought to skip school (perhaps her evidence supports it), then it follows from the wide-scope Self-Reliance Principle that she ought to skip school. This solution implausibly entails the impossibility of having false but *justified* normative beliefs.¹⁵ And if she ought not to abandon her intention to skip school (perhaps she has persuasive evidence that she ought to skip school), then it follows from the wide-scope Subjective Instrumental Principle that she ought to feign illness. This would entail the impossibility of doing the wrong thing *justifiably* in successfully pursuing one's ends.

Efforts to solve these detaching problems without appeal to more than one sense of 'ought' apparently fail, and others propose a solution appealing to just *two* senses of 'ought': an *objective* 'ought' (of having most reason) and a *subjective* 'ought' (of what is rational).¹⁶ The 'ought' claims in the consequents detach, but the contradictions aren't genuine. If Jorja believes she objectively ought to skip school, then she *subjectively* ought to skip school, which is consistent with her belief being false. The Self-Reliance Principle would not therefore allow our normative beliefs about what we ought objectively to do to bootstrap themselves to infallibility. Likewise, if Jorja wrongly intends to skip school, then she subjectively or *rationaly* ought to attempt the believed means—which doesn't contradict the proposition that she objectively ought not to do so. The Subjective Instrumental Principle would thereby be innocent of allowing our intentions to make their own successful pursuit correct.

This solution, and the distinction between subjective and objective 'ought's, may be on the right track. But any hope that we can adequately accommodate these normative

¹⁴ Another case involves the principle that 'ought' implies 'can' (Greenspan 1975; Schroeder 2004; Kolodny 2005). If Jorja psychologically cannot rid herself of the false belief that she ought to skip school, the wide-scope Self-Reliance Principle together with this Ought Implies Can Principle tell us that she ought to skip school, as the only way she can make the conditional true. This implausibly suggests that normative beliefs that are psychologically inescapable can never be false. Similarly, the Subjective Instrumental Principle would tell us that the successful pursuit of inescapable intentions is never wrongful.

¹⁵ It could be objected that Jorja's being *justified* in her false belief is not incompatible with its being the case that she ought not to believe it. This may be true, but only if we accept a distinction between what she subjectively ought to believe and what she objectively ought to believe (Schroeder 2009). Such a rescue of the wide-scoping strategy would scuttle its claim to parsimony.

¹⁶ Ewing (1953); Kolodny (2005); Schroeder (2009). For the distinction, see also Sidgwick (1907, p. 207).

principles by a distinction between just two senses of ‘ought’ is naïve.¹⁷ The problems here can be approached by asking for the meaning of this rational, ‘subjective ought’ and its relation to the ‘objective ought’. The canonical view, which aims both to limit the necessary senses of ‘ought’ to two and to make their relation transparent, is that what you subjectively or rationally ought to do just is what you objectively ought to do were your beliefs about your reasons true.¹⁸ This may seem to fit nicely with the Self-Reliance Principle. If (i) the subjective ‘ought’ concerns what is rational, (ii) what is rational is what one believes one has most reason to do, and (iii) what one objectively ought to do is just what one has most reason to do, then it follows that if Jorja believes that she ought (objectively) to skip school, then she ought (subjectively) to skip school. But this treatment of the principle faces at least two problems.

The first problem concerns the principle’s normativity. If ‘subjectively ought’ simply means ‘objectively ought if one’s beliefs are true’, then we may doubt the normativity of the principle wherever those beliefs are false. The Self-Reliance Principle itself would not tell us that anything is wrong with us if we do not act as we mistakenly believe we ought. Kolodny accordingly suggests that these principles of rationality have only ‘apparent’ normativity (2005, p. 513); they merely concern the appearance from an agent’s own perspective of what she has reason to do. While I agree that rational requirements do not entail practical reasons, this interpretation sits uncomfortably with the *evaluative* intuition, which Kolodny shares, that the person who fails to (try to) act as she mistakenly believes she ought is in some way ‘defective’ or ‘malfunctioning’. If we think such a failure is a genuine malfunction in her will, we should be able to say that in some genuinely normative sense she *ought* not to act that way—meaning more than merely that this is how it falsely appears from her perspective—and this seems precisely what the Self-Reliance Principle is supposed to capture.

The second problem arises from the plausible supposition that there are different kinds of ‘ought’ beliefs. If Jorja believes that *all things considered* she ought to skip school, then it is plausible that she ought ‘rationally’ to skip school. But what if she believes rather that *morally* she ought to skip school? Like many, I believe that it is not necessarily irrational to fail to comply with one’s moral beliefs. The problem is not that the Self-Reliance Principle is false, but that it requires even greater proliferation of senses for ‘ought’. For surely there is a sense in which, if Jorja believes she morally ought to skip school, she ought to skip school. This isn’t (simply) a rational sense; it rather seems to be a kind of moral sense. Jorja is not *morally* as she ought to be

¹⁷ Besides what follows, we need to distinguish between what an agent ought to do relative to (a) her nonnormative beliefs, and (b) her beliefs about what she ought objectively to do. Additionally, there is still a form of bootstrapping problem here. Plausibly, if an agent believes that she *subjectively* ought to do *A*, then she ought in some sense to do *A*. Since it seems that we can be mistaken about what we subjectively or rationally ought to do, this other ‘ought’ cannot have the same subjective sense. We may be able to generate an infinite regress of senses of ‘ought’ this way.

¹⁸ For views in this neighborhood, see Scanlon (1998, pp. 25–30); Davidson (2004); Kolodny (2005, p. 557); Parfit (1997), *forthcoming*. Schroeder (2009) accepts this account provisionally, before proposing an alternative that defines subjective ‘ought’ in terms of subjective reasons. My CP analysis of ‘ought’ below may be compatible with Schroeder’s once reasons are defined in compatible terms (See Finlay 2006). For problems with the canonical view, see Ross (2006).

if she does not follow her conscience. This also cannot be the same moral sense of ‘ought’ that figures in her belief, on pain of moral bootstrapping and contradiction. We apparently need to distinguish between an objective and a subjective *moral* ‘ought’, and parallel arguments suggest a need to distinguish between objective and subjective *prudential* ‘ought’, etc.

The Subjective Instrumental Principle presents a greater difficulty for anything resembling the canonical view. Jorja’s *intending* to skip school seems not necessarily to involve any normative beliefs. So the truth of the antecedent (Jorja intends to skip school and believes feigning illness is the necessary means) doesn’t seem to entail that Jorja has any beliefs such that were they true, she would have most reason (objectively ought) to skip school or take the necessary means. If weakness of will is possible, Jorja can intend to skip school while believing that she has most reason to act otherwise. The disambiguated Subjective Instrumental Principle would then be false, and there would not necessarily be anything defective about an agent who failed to take means she believed necessary for her intended end.¹⁹

The instrumental ‘ought’ resists accommodation by a simple distinction between the objective and subjective ‘ought’s. This prompts a fourth strategy—the final one discussed here—which accounts for the Subjective Instrumental Principle rather by appealing to a further kind of normative ‘ought’: the *epistemic* ‘ought’ of *having sufficient evidence*, the ‘ought’ governing belief. ‘Cognitivists’ about practical reason maintain that intentions to do *A* are in part beliefs that one will do *A*.²⁰ Setiya (2007) accordingly argues for an epistemic reading of the Subjective Instrumental Principle. The crux of the strategy is that if (e.g.) Jorja intends to skip school and believes that feigning illness is the necessary means, then in virtue of the beliefs constitutive of intention, she fails to conform to the laws of theoretical reason if she fails to intend to feign illness. The detaching problem is avoided, because we can safely detach this epistemic ‘ought’ without contradicting any genuinely practical ‘ought’.

This strategy raises many issues, but I shall merely observe that it depends on the narrow formulation of the Principle in terms of *necessary* means. It is hard to see how it could succeed in accommodating the intuitive broader principle that an agent ought to take what she believes to be the *best* means to her intended end.²¹ For it is compatible with *A*’s being the best means that other actions are means too. If other means are available, then it isn’t true that the agent’s believing that she will achieve her end and that doing *A* is the best means rationally requires her to believe that she will do *A* (and thereby intend it).²² The cognitivist strategy looks unpromising.

¹⁹ Kolodny (2005, p. 559) acknowledges this problem for his solution. See also Broome (2002, p. 99); Setiya (2007). Schroeder (2009) proposes that intending an end really does entail believing that there is most reason to bring it about. He confesses that while he finds it hard to believe, he cannot see how else to preserve the Subjective Instrumental Principle.

²⁰ See Harman (1976); Velleman (2000); Wallace (2001), and criticism in Ross (forthcoming).

²¹ This point is also made in Broome (2002, p. 109), and Schroeder (2009).

²² Setiya acknowledges the problem in a closing footnote, and expresses optimism that it can be solved. I see no basis for this optimism. One might try appealing to Davidsonian principles of charity, claiming that we must assume agents to be substantively rational and thereby can expect them to choose the best means. But this illicitly rests on something that the cognitivist has denied himself: the claim that there are practical reasons to choose the best means to one’s end. The cognitivist seems committed to saying that if an agent

A further obstacle to any parsimonious treatment of ‘ought’ emerges from the instrumental cases. Consider the

Objective Instrumental Principle: If an agent *S* intends to achieve an end *E*, and doing *M* is the necessary means to achieving *E*, then *S* ought to do *M*.

This principle concerns the normative relationship between an agent and the *genuinely* necessary means to her intended ends, whatever she may believe.²³ If Jorja intends to skip school, and feigning illness is the necessary means, then (in order to achieve her end) she ought to feign illness, even if she falsely believes that she can achieve her end by feigning tiredness instead. This *objective instrumental* ‘ought’ is clearly distinct from the objective ‘ought’ of having most reason. That Jorja ought (in this sense) to feign illness is compatible with her having most reason not to skip school and therefore not to take the necessary means, and so the detaching problem rears its head; if Jorja wrongly intends to skip school, then it is and is not the case that she ought to feign illness. This ‘ought’ must also be distinct from (i) the subjective ‘ought’, and (ii) the epistemic ‘ought’. Neither failing to take the necessary means to your end nor failing to believe that you will perform some particular action as a means is *rationaly* defective if you are not aware of the necessity of that means. If Jorja incorrectly believes that feigning tiredness is the necessary means to skipping school, it would rather be irrational for her either to attempt the actual means of feigning illness, or to believe both that she will instead feign illness and that she will achieve her end.

We apparently need *yet another* sense of ‘ought’ to accommodate the Objective Instrumental Principle. I have encountered resistance to the suggestion that this normative principle genuinely obtains. But perhaps no other kind of ‘ought’ claim is more common in ordinary discourse than the ‘hypothetical imperatives’ which take this form. The objective instrumental ‘ought’ has a much clearer normative or advisory use than the subjective, rational ‘ought’. The usual situation in which one tells somebody, ‘If you intend/want *E*, then you ought to do *A*’ is precisely when he does *not* already appreciate that *A* is the necessary or best means to achieving *E*.²⁴ We are regularly guided by such information. Rejection of the principle therefore appears question-begging, based on a prior restriction of the senses of ‘ought’ to a subset of those canvassed above.

To sum up: in seeking to accommodate just these three principles and avoid the detaching problem we have encountered at least the following distinctions in senses of ‘ought’: objective vs. subjective, practical vs. theoretical, instrumental vs. non-instrumental, and also the differentiation between the many different domains of the practical ‘ought’ (moral, prudential, etc.) Some of these distinctions cut against each other, so we’ve seen (e.g.) a need for an objective instrumental ‘ought’ and a subjective instrumental ‘ought’. This proliferation of senses threatens confusion in normative

Footnote 22 continued

knows that she is going to choose the inferior of two means, then the Subjective Instrumental Principle tells us that this is what she ought to do.

²³ Similar, familiar principles substitute what an agent *wants* for what she intends, and the *best* means for the necessary means.

²⁴ Finlay (forthcoming a).

philosophy, and we have reason to seek a semantic theory of ‘ought’ that accommodates these principles and helps make sense of the bewildering variety of senses. Fortunately, when we look to standard views in linguistics and deontic logic we find the makings of a solution.

3 The standard semantics for ‘ought’

Among linguists, there is consensus that ‘ought’ is semantically modal, like the other deontic terms ‘must’, ‘have to’, ‘may’, ‘could’, etc. This might appear to make our difficulties worse, because it draws our attention to the fact that like other modals, ‘ought’ has also nonnormative uses, which are commonly called ‘predictive’ or ‘epistemic’.²⁵ In the ethics literature acceptance of genuine semantic ambiguity is commonplace here, even among those most opposed to ambiguity in normative ‘ought’.

The standard treatments of modals find semantic uniformity, however.²⁶ Verbs like ‘must’ and ‘may’ have many shades of surface meaning because they are semantically incomplete, with places for arguments commonly supplied by context. The semantic value of a modal is an invariant operation that it performs on these arguments. Following Kratzer (1981), modals are attributed two argument places in addition to their complements: a ‘modal base’ B , and an ‘ordering source’ O ; their logical form is thus $must_{B,O}(p)$, $ought_{B,O}(p)$, etc. Modals quantify over possible worlds; the modal base B is a set of propositions that restricts the domain of quantification (to the worlds consistent with B), while the ordering source O is some index that yields a particular evaluative ordering of B -worlds. ‘May’ is a possibility modal, and hence an existential quantifier; $may(p)$ is interpreted as meaning that p is the case in some B -worlds (for some salient B) ranked as best by some salient O . ‘Must’ and ‘ought’ are treated as necessity modals and hence universal quantifiers; $must(p)$ and $ought(p)$ are interpreted as meaning that p is the case in all B -worlds ranked as best by O . (The standard semantics controversially assumes, as I will, that ‘ought’ and similar terms always operate on propositions.)²⁷

This approach offers resources for resolving the issues encountered in Sect. 2. It is flexible enough to generate an indefinitely large number of distinct surface senses for

²⁵ Philosophers generally use ‘epistemic’ for the ‘ought’ of belief, while linguists use it for this nonnormative ‘ought’.

²⁶ The seminal texts here are Sloman (1970); Kratzer (1977, 1981).

²⁷ ‘Ought’ is commonly held to have a *relational* sense (‘ought to do’, versus ‘ought to be’) expressing a relation between an agent and an action; e.g. Harman (1973); Vermazen (1977); Geach (1981); Schroeder (ms). This is denied also by Chisholm (1964); Williams (1981); Broome (2002); Wedgwood (2006). Briefly: (1) mere grammatical differences prove nothing, because English commonly requires a grammatical subject even for verbs that have no logical subject. ‘The meeting ought to start at noon’ has the same grammatical form as the supposedly relational ought-sentences, but clearly means that it ought to be the case that the meeting starts at noon (I take the example and point from Schroeder ms). Note that even ‘it ought to be the case...’ has a grammatical subject: ‘it’. (2) The *normative* distinctiveness of the supposedly relational ought-sentences might be explained by an implicit *agency* operator; e.g. ‘Jorja ought to feign illness’ could be elliptical for ‘It ought to be that Jorja sees to it that she feigns illness’ (Belnap and Horty 1996—although Horty gives up the reductive claim in his 2001, for reasons peculiar to his own semantics; see also Sloman 1970; Williams 1981; Broome 1999; for objections see Schroeder ms).

‘ought’, by varying the implicit arguments supplied to the common modal operator. It also offers a straightforward solution to the detachment problem, and an explanation for why normative principles sometimes don’t licence detachment by modus ponens. *If*-clauses conditionalizing modal claims are generally interpreted as functioning to add propositions to the modal base, further restricting the relevant domain of worlds.²⁸ These antecedents therefore supply part of the *meaning* of their consequents, rather than just placing conditions on their *truth*. For example,

(1) If Jorja feigns illness, then she might skip school,

can be parsed as claiming that Jorja *might-if-she-feigns-illness* skip school; its logical form is not $A \rightarrow \Diamond B$ but rather $\Diamond_A B$. The consequent therefore does not itself express a complete proposition, and cannot be detached from the antecedent. To illustrate, consider that one might be able to claim truthfully both (1) and

(2) If Jorja’s parents are wise to her ruses, then she cannot possibly skip school.

Both conditionals can be true together, and both their antecedents too: Jorja feigns illness, but her parents are indeed wise to this ruse. But if we could detach the consequents, we would have a contradiction; Jorja might skip school but she cannot possibly.

Unfortunately this standard semantics has some shortcomings due to the *ordering source* mechanism,²⁹ particularly for ‘ought’;

- (i) There is a clear difference between ‘ought’ and ‘must’, which the standard semantics fails to capture. ‘Ought’ is weaker than ‘must’ (e.g. ‘You don’t have to, but you ought to’) and stronger than ‘may’. It is widely agreed that *must(p)* means that *p* is the only possibility, while *ought(p)* means that *p* is merely the *best* of the possibilities.³⁰ Yet the standard semantics analyzes both as meaning that *p* is the case in all the best worlds.³¹
- (ii) It struggles to accommodate the *nonnormative* ‘ought’. Here also ‘ought’ is intuitively weaker than ‘must’; it is commonly and intuitively interpreted as meaning roughly *probably*.³² The general strategy for accommodating nonnormative modals is to postulate an *empty* ordering source (so that the nonnormative

²⁸ Kratzer (1981); Sæbø (2001, p. 428); Broome (2002, p. 105).

²⁹ The operation of the ordering source also raises questions. It is supposed to accommodate a variety of kinds of index, including intentions and desires, value standards (e.g. what is good for *x*), ideals, rules, laws, and duties. What operation could systematically yield rankings of worlds as ‘best’ and ‘better’ relative to these various items? On Kratzer’s model (1981, p. 47) desires, ideals, rules, etc. supply a set *O* of (sometimes inconsistent) propositions, and a world w_1 is ranked as better than a world w_2 just in case w_2 is consistent with at most a proper subset of the part of *O* with which w_1 is consistent. This is a crude model; it doesn’t allow some ideals to take priority over others, and it doesn’t yield a ranking of two worlds if each has unique members.

³⁰ Sloman (1970).

³¹ McNamara (2006). Solutions are proposed in more recent work; perhaps the most sophisticated is von Fintel and Iatridou (2008), who suggest that ‘ought’ takes *two* ordering source arguments. In lieu of a careful comparison, note that their view (i) is thereby more complex than the one I propose below; and (ii) doesn’t as easily accommodate the nonnormative ‘ought’. Also see footnote 47.

³² For example Sloman (1970); Harman (1973); Wheeler (1974); Nordlinger and Traugott (1997, p. 298); Wedgwood (2006); Thomson (2007).

‘must’ and ‘may’ quantify over *all B*-worlds without restriction from *O*). But because ‘ought’ is weaker than ‘must’, it apparently quantifies over a restricted set of *B*-worlds, and therefore this strategy isn’t available. Consequently, ‘ought’ is supposed to require a nonempty ordering source.³³ A nonnormative ‘ought’ must therefore, anomalously and counterintuitively, select what is ‘best’ relative to some ordering source or other. It is hard to see, however, that when we say that something ‘ought’ (probabilistically) to be the case, we are saying that it is the *best* in some salient sense.³⁴ Some linguists have accordingly endorsed the view that normative and nonnormative ‘ought’s require distinct, even if related, semantic analyses.³⁵

- (iii) It has difficulty with *instrumental* ‘ought’ sentences. (I examine this problem in the next section).
- (iv) For those coming from ethical philosophy, this analysis is disappointing. We might have hoped that discovering the meaning of ‘ought’ would teach us something about the mysterious nature of normativity itself. But the standard semantics accommodates the normativity of ‘ought’, ‘must’, etc. simply by appealing to further assorted normative concepts (‘duty’, ‘best’, ‘good’, ‘rule’) which are no better understood—and may even need to be explicated using terms like ‘ought’—leaving the normativity of ‘ought’ mysterious.

It may be that issues (i)–(iii) can be resolved, and that (iv) asks for too much. But they provide motivation for seeking a better theory. I will now argue that there is an alternative that is superior in all the following ways. (a) It is more unifying, accommodating without difficulty the normative and nonnormative senses of ‘ought’ with a single semantics, in a way that exactly parallels the relation between the normative and nonnormative senses of ‘must’, etc. (b) It offers a simpler semantic model by eliminating the problematic ordering source argument for ‘ought’ and other normative modals. (c) It analyzes ‘ought’ without appeal to other, unexplained normative concepts, giving a determinate and systematic account of how the different senses of ‘best’ are generated—offering an enlightening reduction of normative claims and an explanation of the nature of normativity. (d) Finally, as shown in Sect. 5, it straightforwardly accommodates the three troublesome principles examined in Sect. 2.

4 An alternative semantics for ‘ought’

A good method for finding a better theory is to start by trying to capture the data that fits the extant theories least well. We’ve seen that some of the greatest difficulties for ‘ought’ arise from broadly instrumental cases. This is no less true for the standard semantics. Consider

³³ For example Sæbø (2001, p. 434).

³⁴ Some attempts have been made to explain how being probable is a way of being best; for example, that it is (a) the best thing to believe, or (b) what is best supported by the evidence, or (c) what happens if the world keeps its promises to us. (Sloman 1970; see also von Fintel and Iatridou 2005; Wedgwood 2006; Thomson 2007, p. 263n).

³⁵ Nordlinger and Traugott (1997, p. 298n).

(3) If Jorja wants to skip school, then she ought to feign illness.

On Kratzer's model, the antecedent's function is to add the proposition that Jorja wants to skip school to the contextually-supplied modal base. But the sentence would then seem likely false; plausibly the best worlds relative to the salient ordering source (presumably consisting in what the speaker values) are *not* worlds in which Jorja feigns illness, even if we are restricted to worlds in which Jorja has this desire. Others suggest more plausibly that in sentences like (3), the antecedent rather modifies the contextual ordering source, so the modal addresses what is best relative to an index including Jorja's desire.³⁶ This view also has a problem; the (contextually supplied) ordering source may reasonably be expected to contain other elements, like the value of education or Jorja's desire to pass her tests, and so merely adding Jorja's desire to skip school to *O* won't be sufficient to guarantee the truth of the sentence.³⁷

There is an analysis that avoids these problems. I'll start with 'must', and the following plausible paraphrase (modulo the modal substitution) of (3):

(4) If Jorja is (going) to skip school, then she must feign illness.

Many philosophers and linguists have endorsed such a paraphrase, and have interpreted it as claiming that Jorja only skips school if she first feigns illness.³⁸ Intuitively, the 'must' in (4) is instrumentally normative, and the sentence has the characteristic markers of normativity; it has a clear use in guiding choice and behaviour, and is not falsified if the event that 'must' occur fails actually to occur. Despite its apparent normativity, however, there is no role for an ordering source to play. The sentence seems to reduce to an ordinary claim of conditional necessity.

I propose the following analysis. 'Must' here takes a preliminary modal base B_0 provided by context (consisting roughly of what is known about the state of the world at the time when Jorja is considering trying to skip school, together with such regularities as the laws of physics and psychology). The antecedent modifies this base, yielding an ultimate modal base B , by adding the proposition that Jorja is going to skip school. The consequent ('she must feign illness') then says that in every B -world Jorja feigns illness.³⁹ I call such sentences *end-relational*, using 'end' as a term of art for any proposition expressing a potential future event, irrespective of whether it is actually anyone's goal. In contrast to the standard semantics, this approach takes the nonnormative 'must' as basic, with the logical form $must_B(p)$. We simply eliminate the argument-place for an ordering source; the instrumental normativity of the modal derives rather from conditionalizing the necessity on the obtaining of some hypothetical future event.

³⁶ See Sloman (1970); Sæbø (2001).

³⁷ Huitink (2005) departs from the standard model by suggesting the antecedent *determines* rather than merely modifies the ordering source. As von Stechow et al. (2006) observe, this trivializes the notion of an ordering source, since the same results can be obtained by collapsing it into the modal base, as I suggest below.

³⁸ Bech (1955); von Wright (1963, p. 10f); Sæbø (2001); Broome (2001); von Stechow et al. (2006); von Fintel and Iatridou (2005); Finlay (2009).

³⁹ See also von Stechow et al. (2006).

Is this end-relational analysis plausible? A couple of issues need examination. First, (4) seems to contain elements omitted by the analysis. Many claims of the form *if p then necessarily q* are not normative; for example, ‘If Jorja is skipping school, then she must be breathing.’ There are certainly some other necessary conditions for a normative reading. One is that the sentence must make clear that the consequent identifies a ‘means’ or metaphysically prior condition for the end. This requirement provides an innocent explanation for why an infinitive (e.g. ‘to skip’) is necessary in the antecedent in order to get an instrumental reading; metaphysical priority is conveyed grammatically by temporal priority, but in a normal conditional (e.g. ‘If Jorja skips school, then she feigns illness,’) the consequent is interpreted rather as temporally subsequent to the antecedent. The infinitive is therefore plausibly a device for forcing the antecedent to be interpreted as temporally and metaphysically subsequent to the consequent.⁴⁰ Another condition arguably necessary for capturing the *practical significance* of normative claims is that the end is valued (by the speaker, perhaps, or the agent). This might be pragmatic rather than semantic—i.e. a matter of use rather than content—in which case being an intuitively normative sentence may itself be in part pragmatic. There may be further conditions.⁴¹

Second, there is an issue that will prove important later. If this strategy is to succeed in accommodating instrumental normative claims, it has to factor out the influence of an agent’s psychological dispositions to choose certain actions or means over others. If Jorja could skip school by feigning either illness or tiredness, but given her psychology it is impossible that she chooses to feign tiredness, then the theory seems to predict falsely that we may say, ‘If Jorja is going to skip school, then she must feign illness.’ To avoid this, the theory must suppose that the modal base is insensitive to Jorja’s dispositions, and includes worlds in which she chooses courses of action that she would not actually choose. This move does not place the account at any disadvantage against the standard semantics, since Kratzer notes that she has to make it too (1981, p. 53).

We have a defensible analysis of an instrumental ‘must’ that is simpler than the standard semantics. Can we extend it to all other sentences involving a normative ‘must’? We can if every normative ‘must’ appeals at least covertly to some end or ideal. So claims an end-relational (ER) theory of normativity, which I have elsewhere argued is independently well-motivated.⁴² So, for example, the theoretical ‘must’ assumes something like *if you are going to believe what is true*, and the moral ‘must’ arguably assumes (e.g.) *if you are going to maximize utility/treat humanity always as an end and never only as a means*. All normative necessity would therefore reduce to there being only a single way of bringing about some salient end. The different kinds of normative ‘must’ result from restricting the base by different ends (moral ends, prudential ends,

⁴⁰ I explore this in detail in Finlay (forthcoming b). Also see Sæbø (2001); von Stechow et al. (2006).

⁴¹ A normative reading also requires the modal to be tensed as prior to the end (Finlay forthcoming b). Contrast ‘If Jorja wanted to skip school, then she *has* to have feigned illness’ (epistemic) and ‘If Jorja wanted to skip school, then she *had* to feign illness’ (normative). This would partially explain why tenseless modal operators (like ‘necessarily’, ‘possibly’, and ‘probably’) don’t allow normative readings. (‘Ought’ and ‘must’ are defectively tensed, lacking past tenses, but not tenseless).

⁴² See Finlay (2004, 2006, 2009); also Foot (1972). Arguably Kant himself took this view; the ‘hypothetical’ condition on moral imperatives can be suppressed because every rational agent necessarily aims to satisfy it.

chess ends, etc.) The end-relational theory achieves the same flexibility as the standard semantics with fewer moving—and no mysterious—parts.

I propose a semantics for ‘ought’ that follows exactly this blueprint. Instead of starting with the normative ‘ought’ as basic and trying to analyze the nonnormative ‘ought’ into it (i.e. reduce ‘probably’ to ‘best’), we start with the nonnormative ‘ought’ as basic.⁴³ As others have noted, ‘ought’ seems to refer essentially to implicit classes of relevant alternatives,⁴⁴ so I propose a *comparative probabilistic* (CP) theory of this basic ‘ought’. $Ought_{B,R}(p)$ says that p is more likely than any member of some salient class of alternatives R , given B . I would prefer to be neutral about the interpretation of probability, but for the sake of exposition we can think of probability as a proportion of a finite, non-empty possibility-space.⁴⁵ I define the probability of a proposition p relative to a base B to be the proportion of B -worlds that are worlds in which p obtains (B_p -worlds):

$$\text{pr}_B(p) = \frac{|B_p|}{|B|}$$

Now apply the end-relational theory, and consider the following sentence:

(5) If Jorja is going to skip school, then she ought to feign illness.

The theory interprets this as: given that Jorja will skip school (q), it is more likely that she first feigns illness (p) than any relevant alternative r ;

$$\text{pr}(p|q) > \text{pr}(r|q), \quad \text{for all } r \in R$$

This may seem an implausible analysis of (5); how can ‘most likely’ be equivalent to ‘best’? Jorja’s being *most likely* to feign illness seems quite different from its being *best* that she feigns illness. It becomes less implausible, however, once we faithfully follow the blueprint established above for ‘must’. In order to get the right results for ‘must’, we followed the standard semantics in interpreting the preliminary modal base (i.e. prior to restricting it to the worlds in which the end eventuates) as being insensitive to the agent’s dispositions to either pursue or neglect any of the relevant potential means; i.e. the disposition of any member of $\{p, R\}$ to be true. Therefore the base should involve an assignment of equal initial probability to the choice of every relevant potential means, by the principle of indifference. I will call this assumption of equal initial probability *symmetry of choice*. Given the subsequent base, i.e. that

⁴³ Wheeler (1974) proposes that what A ‘ought’ to do is what A probably does, given that A is a member of its kind (person, believer, etc.)

⁴⁴ Sloman (1970); Jackson (1985).

⁴⁵ To avoid vicious circularity, I must reject at least subjective interpretations of probability couched in terms of what *ought* to be believed. To motivate my model intuitively, consider the probability of the proposition that a certain child has red hair, given only that (i) each parent has one gene for red hair and one gene for fair hair, and (ii) that fair hair is the dominant trait. The possible outcomes are RR, RF, FR, FF. In only one of these four ‘worlds’ will the proposition be true, hence the probability is 0.25. These look like implausible candidates for possible worlds, but we might try a model that quantifies rather over relevant world-types, identifying RR as a world-type and applying the principle of indifference to the different types.

we have symmetry of choice and the end eventuates, the means most likely chosen is the means on which the end has the greatest likelihood of eventuating. This follows by an application of Bayes' Theorem, but can be intuitively grasped independently of any familiarity with Bayesianism.⁴⁶ Imagine that in the morning before school I know that Jorja hopes to skip school, and I also know that there are only two means to doing so: feigning illness, to which I assign a probability of success of 0.9, and feigning tiredness, to which I assign a probability of success of 0.1. I have no idea which she will choose, so (by the principle of indifference) relative to what I know, each choice has a probability of 0.5. Now imagine that in the evening I discover that Jorja somehow managed to skip school. It would be very natural for me to say, 'She probably feigned illness, then' (or: 'She ought to have feigned illness, then').

This CP/ER theory of 'ought' implies that the 'best' means relative to a particular end is the most *reliable* means; i.e. the means on the choice of which the end is more likely than it is on any of the alternative means. This is equivalent to the means that is most likely chosen, given success and symmetry of choice; the fact of success itself raises the probability that a more reliable means was selected. It is plausible that the best means relative to a particular end is the most reliable means. Admittedly, we often evaluate means on other grounds, such as cost, pleasantness, moral permissibility, safety, etc. However these are plausibly evaluations of means relative to ends *other* than the ends to which they are means.⁴⁷ In ordinary conversation speakers are concerned with a variety of ends; this introduces complexities into our use of 'ought' that I will not attempt to address in this paper, but that would need to be accommodated for this account to be fully successful.

This application of the end-relational theory to the comparative probabilistic semantics for 'ought' is therefore a reductive account of normative 'ought' that yields intuitively correct results for instrumental ought-claims. Suppose further, as before, that every normative 'ought' is at least implicitly end-relational; every normative 'ought' would thereby involve the greater reliability of some course of action for promoting some salient end over any relevant alternative. On this analysis, 'ought' takes three arguments—its complement p , its modal base B , and a set of relevant alternatives R . Its logical form is $ought_{B,R}(p)$, and its meaning is that p is more likely given B than any alternative in R . Normative senses of 'ought' arise from restricting the base to worlds where some future end/event occurs. Unlike the standard semantics there

⁴⁶ This analysis has been accused of implausibly attributing a grasp of Bayes' Theorem to ordinary users of 'ought'. But the point is simpler than Bayes' Theorem (since it only involves ranking probabilities, not calculating them). Furthermore, it isn't true that people are ignorant about how to update their credences in light of evidence until they are taught Bayes' Theorem. We learn how to do it intuitively at a young age—although even adults typically do it poorly.

⁴⁷ Contrast von Fintel and Iatridou (2005), who propose that instrumental 'ought' picks out the best relative to an agent's *other* goals. Huitink (2005) argues, I think correctly, that this generates the wrong results. To motivate my proposal, suppose you have to choose between three doors; door A is certain death, door B has 0.7 probability of death, and door C has 0.5 probability of death. It seems correct to say that if you want to live (and nothing else matters to you) then you *ought* to choose door C but you don't *have* to, since you *could* (albeit irrationally) choose B. Take away door B, however, and it seems correct to say that if you want to live, then you *have* to choose C.

is no argument place for an ordering source⁴⁸; different senses of *best* are generated by supplying different bases (ends and circumstances) to the common *comparative probability* function.

Compared to the standard semantics, this CP/ER approach provides a simpler, clearer, and yet more unifying account, with greater reductive and explanatory power. It offers an explanatory reduction of normative ‘ought’ propositions about the actual world as simply propositions about what is most likely the case in nearby worlds where there is symmetry of choice and certain ends eventuate. (As with ‘must’ there will be additional necessary conditions, since probabilistic claims are typically not intuitively normative).⁴⁹ Judgments about what we ought to do if we are to curb global warming, for example, are simply judgments about what we would probably do were we subsequently to prove successful in curbing global warming. Joseph Butler famously wrote that ‘probability is the very guide of life’; since the same may be said of ‘ought’, it should not be too remarkable if they turn out to be the same thing.

5 A solution to the detaching problem

I have argued that the CP/ER analysis is simpler, more unifying, and more explanatory than the standard semantics. In this section I apply it to the three puzzling principles discussed in Sect. 2, show that it makes sense of them and provides straightforward solutions to our puzzles. A fully satisfactory treatment of these principles will meet the following criteria. (1) It avoids the detaching problem (bootstrapping and contradiction). (2) It offers a sufficiently *transparent* interpretation, in that ordinary language users will be able to recognize utterances of those sentences as expressing the suggested propositions, identifying the required contextual parameters. (3) It preserves the intuitive *truth* of the principles. (4) It captures their normative *force*.

The so-called ‘objective ought’ on this approach concerns what makes some (implicit or explicit) end or ideal comparatively most likely. Moral ‘oughts’ would then address moral ends or ideals, prudential ‘oughts’ address ends of prudence, etc. If we are to start with the simplest case, we ought to address first the Objective Instrumental Principle, which tells us (e.g.) that if Jorja intends to skip school, and feigning illness is in fact the necessary means, then she ought to feign illness.

⁴⁸ Two issues: (1) Is the analysis *simpler*, if it substitutes a relevant alternatives argument-place for the ordering source argument place? It offers a simpler account of normative modals in general, and making ‘ought’ essentially comparative is independently well-motivated and has been proposed elsewhere (Sloman 1970; Jackson 1985). (2) Are there independent reasons for recognizing a place for an ordering source? Kratzer identifies some, but my account has resources to neutralize them. (i) Her ‘main motivation’ (early in her paper) is ‘the necessity to obtain notions of graded possibility’ (1981, p. 53). But the probabilistic model also achieves this. (ii) A major motivation is accommodating ‘human’ necessities where even the best worlds are suboptimal. Kratzer sees this as requiring the modals to be restricted by inconsistent propositions which if placed in the modal base would yield empty sets. The end-relational account, however, provides a very different model for normative claims. Kratzer also utilizes the ordering source to analyze counterfactuals; my account would require a reassessment.

⁴⁹ As with ‘must’, these likely include a relation of metaphysical priority and that the end is valued. See also footnote 41.

The CP theory tells us that the ‘ought’ here is to be interpreted as *probably given* B , for some base B . Since the principle is intuitively normative, the ER theory tells us that B has been restricted at least to the obtaining of some end e ; the ‘ought’ is equivalent to *probably given* B_1 and that it is to be that e , for some circumstantial base B_1 such that $B = \{B_1, e\}$. What might we plausibly identify e and B_1 as? The most natural candidate for a *preliminary* modal base (B_0), given presumable conversational interests, is simply the actual circumstances known to the speaker in which (e.g.) Jorja has to deliberate and act, with symmetry of choice. This preliminary base B_0 is overtly modified by the hypothesized fact m that feigning illness is the necessary means to skipping school, to yield the circumstantial base $B_1 = \{B_0, m\}$. The obvious candidate for the salient end e is then Jorja’s end, explicitly introduced in the antecedent, that she skips school.⁵⁰ The Objective Instrumental Principle then tells us, of Jorja’s situation, that given B (that she will skip school, for which feigning illness is the necessary means, in the known circumstances in which she is to deliberate and act, with symmetry of choice), she more likely feigns illness than any relevant alternative.

This analysis meets the transparency criterion quite well; both the (relevant modification of the actual) circumstances and the relevant end are explicitly introduced in the antecedent. The CP/ER analysis is tailor-made to accommodate the Objective Instrumental Principle. There is, however, a difference between ‘if Jorja is going to skip school’ and ‘if Jorja intends to skip school’. What is the relevance here of Jorja’s *intentions*? The comparative probability of skipping school conditional on taking the necessary means does not depend upon her intentions. However, while the *truth conditions* of the ‘ought’ claim, so interpreted, do not depend upon Jorja’s intentions, its *relevance conditions* do depend on them. In giving instrumental advice we are concerned with the promotion of ends that agents care about. ‘If Jorja intends...’ would then function to introduce the conversational relevance of the ‘ought’ claim. It is akin to ‘biscuit conditionals’: e.g. ‘If you want biscuits, there are some on the table’.⁵¹ We could paraphrase the sentence as, ‘Given that Jorja is going to skip school, *as she may intend*, and that feigning illness is the necessary means to that end, she ought to feign illness.’

This interpretation preserves the *truth* of the Objective Instrumental Principle. The necessary means is always more likely to be performed than any alternative, given that the end eventuates. Further, elimination of the detaching problem is overdetermined. In addition to the point above that modals cannot be detached from their conditions, observe: (i) Even if we could detach the consequent, the ‘ought’ detached from e would be merely probabilistic and no longer normative, since on the ER theory being end-relational is what makes it normative. (ii) Our interest in the conditional is *as* an interesting counterfactual (what would be probable supposing the end will obtain), so even to think about detaching the consequent is to miss the point. (iii) This normative

⁵⁰ An alternative interpretation identifies e as roughly the accomplishment of the agent’s goals/ maximization of her preference satisfaction. This is not as transparent (this end isn’t explicitly introduced), although the antecedent’s focus on the agent’s intentions may be sufficient to introduce it. The antecedent would then have truth-conditional rather than simply relevance dependence on the antecedent’s information about the agent’s intention.

⁵¹ See also von Fintel and Iatridou (2008), and discussion in Dreier (2009), von Fintel and Iatridou (2005).

‘ought’ has a specific sense; it is ‘ought given B_1 and that it is going to be the case that e ’, and cannot contradict any ‘ought’ taking different arguments. (iv) Finally, only the relevance and not the truth conditions of the proposition depend on the agent’s psychological attitudes. In sum, there is no danger of bootstrapping or contradiction problems.

Turning to the *normative force* criterion, the analysis tells us the following about the Principle’s normativity. It offers no guidance about whether to pursue the end, but guides us in case we are seeking the end by telling us how it may most reliably be attained. If Jorja does as it says she ought then she is more likely to achieve her end than if she does not. This is exactly the normative force the Principle should have. All four criteria we identified for a satisfactory analysis are met.

Next in simplicity is the Self-Reliance Principle, which tells Jorja that if she believes she ought to skip school then she ought to skip school. Here ‘ought’ appears twice in the same sentence, which is (in light of bootstrapping issues) strong evidence of a change in the implicit arguments from the antecedent to the consequent clause. In general, such changes will be transparent to an audience when they are due to information contributed by the antecedent clause itself. The antecedent seems to introduce only the (hypothesized) circumstance that Jorja believes she ought to skip school. So a natural hypothesis is that the second, subjective ‘ought’ differs in meaning from the first ‘ought’ simply in having the proposition about Jorja’s normative belief added to its base; i.e. the first ‘ought’ concerns what is most likely given that B_0 and some implicit end e eventuates, and the second ‘ought’ concerns what is most likely given all this *and* that Jorja believes that skipping school is most likely (given that e eventuates, etc.)

This analysis may not be utterly hopeless,⁵² but it looks hard to capture the *truth* in the Self-Reliance Principle this way. An agent’s believing that some action is most likely given the end’s eventuating doesn’t always make the action most likely given the end’s eventuating. There is a better hypothesis, although it involves a slight loss of transparency. Rather than updating the base with the fact *of* the agent believing that she ought to do A , we update it with the fact *which* she believes (just as in the previous case we updated the base not with the fact that she intends the end, but with the end *which* she intends). That is, we add to the base the supposition that she ought to do A . This secures the Principle’s truth, but its transparency and normative force need discussion.

Regarding transparency, why would we mention the agent’s beliefs if they have nothing to do with the truth conditions? I suggest that the antecedent clause, ‘If Jorja believes she ought to skip school,’ is more perspicuously expressed as ‘If, *as she may believe*, Jorja ought to skip school’.⁵³ We invoke the agent’s beliefs for reasons of relevance again; we are addressing probabilities relative to counterfactuals in which things stand as the agent may believe, which are relevant *because* the agent may have

⁵² Given the assumption that people’s normative beliefs are true more often than not, and no other background information, it would follow from Jorja’s having this belief that she probably ought to skip school. However, this assumption is dubious, and the Self-Reliance Principle seems to apply even when disconfirming background information is available.

⁵³ See Dancy (2000, Chap. 6).

these beliefs. This is not *ad hoc*, as we can observe similar devices in other contexts, including our practices of ascribing ‘motivating reasons’. If asked, ‘What was Jorja’s reason for skipping school?’ an appropriate answer may be, ‘Her reason was that there would be a test.’ But if we know that Jorja was mistaken and there was no test, we should rather say, ‘Her reason was that *she believed* that there would be a test.’ Although we report two different facts, it is implausible that Jorja acts for a different reason in the two cases. One solution is that in the former case, we omit mention of an important part of Jorja’s reason, her belief. But the right solution, I believe, is that although in both cases her reason is *that there would be a test*, saying this has a factive connotation and suggests that the speaker believes it too. To cancel this suggestion, we interpose ‘she believes’.⁵⁴ The analysis of the Self-Reliance Principle thus provides sufficient transparency.

This treatment of the Self-Reliance Principle may seem problematic in respect of normative force. First, it looks tautological and therefore *trivial*, telling us that if *S* ought to do *A*, then *S* ought to do *A*. Surely the Principle is more interesting than that. Second, this doesn’t seem to respect the *difference in sense* between the two ‘oughts’ and the distinctive character of the second ‘ought’. The first is an objective ‘ought’ of having most reason (of some kind), while the second *should* be the subjective ‘ought’. Third, this interpretation appears unable to accommodate the *extension* of the Principle’s normativity. It seems to say only that *if* Jorja’s normative belief were correct, *then* she ought to skip school. Like the standard interpretations, it would then say nothing about what she ought to do in case her belief is false. But as we’ve observed, the Principle captures a defect in behaviour even where the normative beliefs are false.⁵⁵

The answer to all three objections is that there *is* a change in sense from the first to the second ‘ought’. The base of the second ‘ought’ is updated with information provided in the antecedent, as the criterion of transparency suggests. This is the information that *S* ‘may believe’, i.e. that it is more likely that she does *A* than any alternative, given that *e*. Whereas the first ‘ought’ means *most likely given that e and B₁*, the second means *most likely given that e, B₁, and that it is most likely that S does A, given that e and B₁*. Formally, when ‘*m*’ represents *S* does *A*, the Self-Reliance Principle says

$$\begin{aligned} & \text{pr}(m|(e \& B_1 \& [\text{pr}(m|(e \& B_1)) > \text{pr}(r|(e \& B_1))], \text{ as } S \text{ believes})) \\ & > \text{pr}(r|(e \& B_1 \& [\text{pr}(m|(e \& B_1)) > \text{pr}(r|(e \& B_1))], \text{ as } S \text{ believes})), \text{ for all } r \in R. \end{aligned}$$

To clarify, applied to Jorja’s case (and omitting the stable preliminary modal base *B₀*), the Self-Reliance Principle tells us that given (a) that—as she may believe—she more likely skips school than any relevant alternative, given that she will miss the test, and (b) that she will indeed miss the test, it is more likely that she skips school than

⁵⁴ See also Dancy (2000); Kolodny (2005); Finlay (2006); Schroeder (2008). The main defender of the opposing view is Michael Smith.

⁵⁵ An alternative CP/ER interpretation postulates a change in the relevant *end* between the first and second ‘ought’s. If the first is ‘ought given that *e* and *B₁*’ then the second is roughly ‘ought given that you act in accordance with your own normative beliefs’. So interpreted the Principle is true, avoids the detaching problem, and has similar normative force to that of the simpler interpretation. Transparency could be accommodated by assuming that if normative advice is prefixed with ‘If you believe that you ought...’, then it gives advice about how to act in light of your beliefs. As advice it is trivial—but so is the Principle.

that she does anything else instead. Because probability conditionals do not detach, this is not a logical truth or tautology (its logical form is not $(A \& B) \rightarrow A$), but it is a necessary truth given laws of probability. The detaching problem is dispatched for the same reasons pointed out for the previous principle.

The Self-Reliance Principle offers no guidance about how to promote any end or which ends to pursue. It tells us that if an agent is not pursuing the action she thinks best for achieving some end, then by her own lights she is acting in a way less than optimal with regard to that end. This approach reduces the distinction between objective and subjective (rational) senses of ‘ought’ to the familiar distinction between probability relative to what is (a) actual, and (b) believed. It might therefore appear to fall prey to the same *extensional* difficulty pressed against the standard interpretation of subjective ‘ought’ in Section 2: that it tells us only that *if* Jorja’s normative belief is true, then she ought normatively to skip school, whereas the Principle captures a kind of failure even when Jorja’s belief is false. This is not so, although the difference is subtle. According to the CP/ER account, normative ‘ought’s about the actual world just are probabilistic ‘ought’s quantifying over counterfactual worlds that are relevantly related to the actual world. The distinctive normative character of any such ‘ought’ is a function of how those counterfactual worlds diverge from our own. A subjectively normative or rational ‘ought’ is one that quantifies over worlds that are how some agent believes the actual world to be. Such an ‘ought’ is therefore itself a distinctive kind of normative ‘ought’, and doesn’t simply mean ‘ought normatively if p ’.

One might question whether the Principle so interpreted has any normative force, since it seems so trivial. This skepticism is reasonable, but it is unproblematic for the analysis, since the Self-Reliance Principle provides advice of doubtful use in any case. ‘Normativity’ has both a broad and a narrow sense, however. In the narrow sense of *providing guidance to deliberation*, the Principle is arguably nonnormative. But in the broader (‘evaluative’) sense of *providing criteria for external criticism*, the Principle retains its normativity, as it should. There is ongoing debate over whether principles of rationality are normative in the narrow sense or merely in the broad sense,⁵⁶ the CP/ER account preserves this controversy as the question of whether the optimality of an act given that one’s beliefs are true can have any guiding function for deliberation independent of the guiding function of one’s belief that the act is optimal. I can imagine one legitimate pragmatic use for the Self-Reliance Principle: to point out to someone who doubts the veracity of her own normative beliefs that ultimately she doesn’t have any other option than to act on their basis, perhaps thereby indicating the speaker’s (and others’) inability or unwillingness to provide advice or information.

The analysis also easily accounts for the difference between ‘rational’, moral, and other varieties of Self-Reliance Principle. Subjective normative ‘ought’s are relativized to ends just as objective normative ‘ought’s are. If Jorja’s objective normative belief is moral (concerns what is likely to promote a moral end), then application of the Principle tells us that she subjectively *morally* ought to skip school. If Jorja is not committed to moral ends, then failure to follow her conscience would not be

⁵⁶ See Raz (2005); Kolodny (2005, p.551f); Scanlon (2003); My view (Finlay forthcoming a) is that irrationality strictly speaking is impossible.

(instrumentally) irrational, but defective rather in a moral sense, a failure to respond to perceived moral reasons.

We now have all the tools needed to explain the Subjective Instrumental Principle, which merely combines the features of the other principles. We read it more perspicuously as, ‘If it is going to be that e (as S may intend), and given that e will not obtain without m (as S may believe), then it ought to be that m .’ Or:

$$\text{pr}(m|(e \text{ [as } S \text{ may intend]} \ \& \ B_1 \ \& \ (e \rightarrow m \text{ [as } S \text{ may believe]}))) \\ > \text{pr}(r|(e \ \& \ B_1 \ \& \ (e \rightarrow m))), \text{ for all } r \in R$$

This shares the transparency of the other two principles; the explicit introduction in the antecedent of the end and the believed necessity of the means informs the sense of the ‘ought’ in the consequent, and the reference to the agent’s intentions and beliefs indicates the relevance of the claim to the agent’s goals and epistemic situation. It is necessarily although not tautologically true; necessarily, Jorja is more likely to feign illness than any alternative, given that she succeeds in skipping school and that her belief about the means is true. The normative ‘ought’ does not detach for the many reasons already noted. The normative force of this principle is also hybrid; like the Self-Reliance Principle it offers no guidance about which ends to pursue, or how to promote those ends. But it tells us that if an agent is not acting in a certain way, then given her own view of the world she is not doing what best promotes her intended ends. There is again room for doubt about whether this principle is normative in the narrow, guiding sense, while it is less controversially normative in the broad, evaluative sense, as is appropriate. Further, it can easily be extended to cases where agents recognize means other than the best means. Formally,

$$\text{pr}(m|(e[\text{as } S \text{ may intend}] \ \& \ B_1 \ \& \ \mathbf{pr}(m|(e \ \& \ B_1)) > \mathbf{pr}(r|(e \ \& \ B_1)) \text{ [as } S \text{ may believe]})) > \text{pr}(r|(e \ \& \ B_1 \ \& \ \mathbf{pr}(m|(e \ \& \ B_1)) > \mathbf{pr}(r|(e \ \& \ B_1)))), \text{ for all } r \in R.$$

I have offered analyses of the three principles, reducing their ‘ought’s to the CP ‘ought’ and solving their detaching problems. Given that the CP/ER account gives a simple, unifying, and reductive analysis of the various senses of ‘ought’, transparently accommodating the three troublesome principles and solving their detaching problems, it is probably the correct account of the meaning of ‘ought’. That is to say, given all these virtues it *ought* to be correct. However as we know, probability conditionals can’t be detached from their conditions. This paper has focused on what the CP/ER approach *can* do, but it has not examined what it *cannot* do. Whether we ought to accept it, all things considered, is the question of whether it is most likely true given what it can *and* what it cannot do.

References

- Bech, G. (1955). *Studien über das Deutsche Verbum Infinitum*. Tübingen: Max Niemeyer Verlag.
- Belnap, N., & Horty, J. (1996). The deliberative stit: A study of action, omission, ability, and obligation. *Journal of Philosophical Logic*, 24, 583–644.

- Bratman, M. (1987). *Intention, plans and practical reason*. Cambridge, MA: Harvard University Press.
- Broome, J. (1999). Normative requirements. *Ratio*, 12(4), 398–419.
- Broome, J. (2001). Normative practical reasoning. *Supplementary Volume-The Aristotelian Society*, 75, 175–193.
- Broome, J. (2002). Practical reasoning. In J. L. Bermudez & A. Millar (Eds.), *Reason and nature: Essays in the theory of rationality*. New York: Oxford University Press.
- Chisholm, R. (1964). The ethics of requirement. *American Philosophical Quarterly*, 1, 147–153.
- Dancy, J. (1977). The logical conscience. *Analysis*, 37, 81–84.
- Dancy, J. (2000). *Practical reality*. New York: Oxford University Press.
- Davidson, D. (2004). Incoherence and irrationality. In D. Davidson (Ed.), *Problems of rationality*. New York: Oxford University Press.
- Dreier, J. (2009). Practical conditionals. In D. Sobel & S. Wall (Eds.), *Reasons for action* (pp. 116–133). Cambridge: Cambridge University Press.
- Ewing, A. C. (1953). *Ethics*. London: English Universities Press.
- Finlay, S. (2004). The conversational practicality of value judgement. *Journal of Ethics*, 8(3), 205–223.
- Finlay, S. (2006). The reasons that matter. *Australasian Journal of Philosophy*, 84(1), 1–20.
- Finlay, S. (2009). Oughts and ends. *Philosophical Studies*, 143, 315–340.
- Finlay, S. (forthcoming a). Against all reason: Skepticism about the instrumental norm. In C. Pigden (Ed.), *Hume, motivation and virtue*. MacMillan.
- Finlay, S. (forthcoming b). Normativity, necessity, and tense. In R. Shafer-Landau (Ed.), *Oxford studies in metaethics 5*. New York: Oxford University Press.
- Foot, P. (1972). Morality as a system of hypothetical imperatives. *Philosophical Review*, 81, 305–316.
- Geach, P. (1981). Whatever happened to deontic logic?. *Philosophia*, 11, 1–12.
- Gensler, H. (1985). Ethical consistency principles. *Philosophical Quarterly*, 35, 156–170.
- Greenspan, P. (1975). Conditional oughts and hypothetical imperatives. *The Journal of Philosophy*, 72(10), 259–276.
- Hampton, J. (1998). *The authority of reason*. Cambridge: Cambridge University Press.
- Harman, G. (1973). Review of Wertheimer. The significance of sense. *Philosophical Review*, 82(2), 235–239.
- Harman, G. (1975). Reasons. *Critica*, 7, 3–18.
- Harman, G. (1976). Practical reasoning. *Review of Metaphysics*, 29, 431–463.
- Hill, T., Jr. (1973). The hypothetical imperative. In T. Hill, Jr. (Ed.), *Dignity and practical reason in Kant's moral theory*. Ithaca, NY: Cornell University Press.
- Horty, J. F. (2001). *Agency and deontic logic*. New York: Oxford University Press.
- Huitink, J. (2005). Analyzing anankastic conditionals. Retrieved 1 Nov 2008 from <http://ncs.ruhosting.nl/sub9/program/huitink.pdf>.
- Jackson, F. (1985). On the semantics and logic of obligation. *Mind*, 94, 177–195.
- Kolodny, N. (2005). Why be rational? *Mind*, 114, 509–563.
- Kolodny, N. (2007). How does coherence matter? *Proceedings of the Aristotelian Society*, 107, 229–263.
- Korsgaard, C. (1997). The normativity of instrumental reason. In G. Cullity & B. Gaut (Eds.), *Ethics and practical reason*. New York: Oxford University Press.
- Kratzer, A. (1977). What *must* and *can* must and can mean. *Linguistics and Philosophy*, 1, 337–355.
- Kratzer, A. (1981). The notional category of modality. In H. J. Eikmeyer & H. Rieser (Eds.), *Words, worlds, and contexts* (pp. 38–74). Berlin: Walter de Gruyter.
- McNamara, P. (2006). Deontic logic. In E. Zalta (Ed.), *Stanford encyclopedia of philosophy*. Retrieved 1 Nov 2008 from <http://plato.stanford.edu>.
- Nordlinger, R., & Traugott, E. (1997). Scope and the development of epistemic modality. *English Language and Linguistics*, 2, 295–317.
- Parfit, D. (1997). Reasons and motivation. *Supplementary Volume-The Aristotelian Society*, 71, 99–130.
- Parfit, D. (forthcoming). *On what matters*.
- Piller, C. (2007). Ewing's problem. *European Journal of Analytic Philosophy*, 3, 43–65.
- Raz, J. (2005). The myth of instrumental rationality. *Journal of Ethics and Social Philosophy*, 1, 2–28.
- Ross, J. (2006). *Acceptance and practical reason*. Rutgers University Ph.D. dissertation.
- Ross, J. (forthcoming). How to be a cognitivist about practical reason. In R. Shafer-Landau (Ed.), *Oxford studies in metaethics 4*. New York: Oxford University Press.

- Sæbø, K. J. (2001). Necessary conditions in a natural language. In C. Féry & W. Sternefeld (Eds.), *Audiatu Vox Sapientiae: A Festschrift for Arnim von Stechow* (pp. 427–449). Berlin: Akademie Verlag.
- Scanlon, T. M. (1998). *What we owe to each other*. Cambridge: Harvard University Press.
- Scanlon, T. M. (2003). Metaphysics and morals. *Proceedings and Addresses of the American Philosophical Association*, 77, 7–22.
- Schroeder, M. (2004). The scope of instrumental reason. *Philosophical Perspectives*, 18, 337–364.
- Schroeder, M. (2008). Having reasons. *Philosophical Studies*, 139, 57–71.
- Schroeder, M. (2009). Means-end coherence, stringency, and subjective reasons. *Philosophical Studies*, 143, 223–248.
- Schroeder, M. (ms). Do oughts take propositions?
- Setiya, K. (2007). Cognitivism about instrumental reason. *Ethics*, 117, 649–673.
- Shafer-Landau, R. (2001). *Moral realism: A defence*. New York: Oxford University Press.
- Sidgwick, H. (1907). *The methods of ethics*. London: MacMillan.
- Sloman, A. (1970). “Ought” and “Better”. *Mind*, 79, 385–394.
- Thomson, J. J. (2007). Normativity. In R. Shafer-Landau (Ed.), *Oxford studies in metaethics 2* (pp. 240–265). New York: Oxford University Press.
- Velleman, J. D. (2000). *The possibility of practical reason*. New York: Oxford University Press.
- Vermazen, B. (1977). The logic of practical “ought” sentences. *Philosophical Studies*, 32, 1–71.
- von Fintel, K., & Iatridou, S. (2005). What to do if you want to go to Harlem: Anankastic conditionals and related matters. Retrieved 1 Jan 2006 from <http://mit.edu/fintel/www/harlem-rutgers.pdf>.
- von Fintel, K., & Iatridou, S. (2008). How to say *Ought* in foreign: The composition of weak necessity modals. In J. Guéron & J. Lecarme (Eds.), *Time and modality (Studies in Natural Language and Linguistic Theory)* (Vol. 75, pp. 115–141). Springer.
- von Stechow, A., Krasikova, S., & Penka, D. (2006). Anankastic conditionals again. Retrieved 1 Nov 2008 from <http://folk.uio.no/torgriso/fs/files/stechowetal.pdf>.
- von Wright, G. H. (1963). *Norm and action*. London: Routledge and Kegan Paul.
- Wallace, R. J. (2001). Normativity commitment and instrumental reason. *Philosophers' Imprint*, 1, 1–26.
- Way, J. (forthcoming). Defending the wide-scope account of instrumental reason. *Philosophical Studies*.
- Wedgwood, R. (2006). The meaning of ought. In R. Shafer-Landau (Ed.), *Oxford studies in metaethics 1* (pp. 127–160). New York: Oxford University Press.
- Wertheimer, R. (1972). *The significance of sense: Meaning, modality, and morality*. Ithaca: Cornell University Press.
- Wheeler, S., III (1974). Inference and the logical “ought”. *Nous*, 8(3), 233–258.
- Williams, B. (1981). “Ought” and moral obligation. Reprinted from *Moral luck*, pp. 114–123 by B. Williams, Ed., Cambridge: Cambridge University Press.