

**CALL #:** [http://sfx.wisconsin.edu/wisc?url\\_ver=Z39.88-2004&ctx\\_ve...](http://sfx.wisconsin.edu/wisc?url_ver=Z39.88-2004&ctx_ve...)  
**LOCATION:** GZM :: EJournals :: Highwire Press Duke University Press

TYPE: Article CC:CCG  
JOURNAL TITLE: <The >Philosophical review  
USER JOURNAL TITLE: Philosophical Review  
GZM CATALOG TITLE: The Philosophical review  
ARTICLE TITLE: "Ought, Agents, and Actions"  
ARTICLE AUTHOR: Mark Schroeder  
VOLUME: 120  
ISSUE:  
MONTH: January  
YEAR: 2011  
PAGES: 1-41  
ISSN: 0031-8108  
OCLC #:  
CROSS REFERENCE ID: [TN:699237][ODYSSEY:129.93.17.103/ILL]  
VERIFIED:

**BORROWER:** LDL :: Main Library  
**PATRON:**



This material may be protected by copyright law (Title 17 U.S. Code)  
System Date/Time: 12/5/2012 12:22:43 PM MST

# Ought, Agents, and Actions

---

*Mark Schroeder*

University of Southern California

## 1.1. Introduction

According to a naive view sometimes apparent in the writings of moral philosophers, ‘ought’ often expresses a relation between *agents* and *actions*—the relation that obtains between an agent and an action when that action is what that agent ought to do. It is not part of this naive view that ‘ought’ always expresses this relation—on the contrary, adherents of the naive view are happy to allow that ‘ought’ also has an *epistemic* sense, on which it means, roughly, that some proposition is likely to be the case, and adherents of the naive view are also typically happy to allow that ‘ought’ also has an *evaluative* sense, on which it means, roughly, that were things ideal, some proposition would be the case.<sup>1</sup> What is important to the naive view is not that these other senses of ‘ought’ do not exist

Special thanks to audiences at the University of California, Irvine, and the University of Nebraska, Lincoln, and particularly to, among those I can remember, Barry Lam, Stephen Finlay, John Broome, Ralph Wedgwood, Jeff Horty, Shyam Nair, Paul Pietroski, Jake Ross, Brian Blackwell, Janice Dowell, Mark van Roojen, John Gibbons, Jeff King, Matthew Chrisman, Maria Nelson, Lewis Powell, Billy Dunaway, Geoff Georgi, Jim Higginbotham, Delia Graff, Sergio Tenenbaum, Amy Challen, Mike McGlone, Ant Eagle, Sari Kisilevsky, Jeff Speaks, Matt King, and the anonymous referees and editors for the *Philosophical Review*.

1. Geach (1982) is the notable case of dissent. Geach believes that ‘ought’ has only the deliberative use on which it expresses a relation between agents and actions (conceived, as I do, as properties) and holds that other apparent uses of ‘ought’ are merely elliptical for deliberative uses. I won’t take this view seriously in this essay. As we’ll see, the evidence for an evaluative sense of ‘ought’ is compelling, so even Geach should grant the existence of such a sense, even if he wishes to go on to analyze it in terms of a deliberative sense.

but rather that they are not exhaustive—for what they leave out is the important *deliberative* sense of ‘ought’, which is the central subject of moral inquiry about what we ought to do and why—and it is this deliberative sense of ‘ought’ that the naive view understands to express a relation between agents and actions.<sup>2</sup>

In contrast, logically and linguistically sophisticated philosophers—with a few notable exceptions<sup>3</sup>—have rejected this naive view. According to a dominant perspective in the interpretation of deontic logic and in linguistic semantics, for example, articulated by Roderick Chisholm (1964) and Bernard Williams (1981) in philosophy and in the dominant paradigm in linguistic semantics as articulated in particular by the most straightforward reading of Kratzer 1977, 1981, there is no argument-place for an agent in any relation expressed by ‘ought’, nor is there any argument-place for an action.<sup>4</sup> According to this view, if Jim ought to jam, that is not because there is a special distinctive deliberative *ought* relation between Jim and jamming; rather, it is because a certain proposition ought to be the case: namely, that Jim jams. The meaning of ‘ought’ is no different, on this view, between ‘Jim ought to jam’ and ‘There ought to be world peace’—in both cases it says merely (roughly) that some proposition would obtain, were things to be ideal. More recently, John Broome (1999, n.d.) and Ralph Wedgwood (2006, 2007) have agreed with the naive view that there is a deliberative sense of ‘ought’ that is distinct from its evaluative sense, and on which it expresses

2. Let me be clear that throughout this essay I will be relying on an intuitive philosopher’s notion of ambiguity and of “senses” of a term, not a more fine-grained linguist’s notion. In particular, I don’t mean to take a stand on any number of more fine-grained questions, including whether the different “senses” of ‘ought’ are like the ‘bank’-‘bank’ ambiguity, the distinction between singular and plural readings of ‘deer’, the difference between being the ‘head’ of a corporation and a physiological ‘head’, the noun ‘strike’ as compared to the verb ‘strike’, or transitive versus intransitive uses of ‘blow up’. If you count only some of these pairs as genuinely “ambiguous,” then please translate my use accordingly. I will loosely call a term ‘ambiguous’ either if it requires different argument structures in different sentences or if it gives rise to distinct truth-conditions.

3. Here I have in mind particularly von Wright (1951), Castañeda (1981), Geach (1982), and ultimately Horty (2001).

4. I say that this is the most natural way of reading Kratzer’s view because her semantics offers no role to be played by an agent. Her semantics *does* appeal, however, to an *ordering source* argument; on some natural ways of developing the view, this ordering source could be set differently for each agent. On this version of the view, it would turn out that agents do play a special role—in special cases when they play a role in setting this ordering source argument. I take it that my central arguments in part 2 are relevant to whether this view could be right, as well.

a relation, one of whose argument-places is for an agent. But Broome and Wedgwood still depart from the naive view in claiming that this relation has no argument-place for an *action*—instead, they claim, it relates agents to arbitrary *propositions*.

It is the aim of this essay to motivate and defend the naive view (although hopefully, without naïveté!) over both the Chisholm/Williams/Kratzer view and against the more similar Broome/Wedgwood alternative. The aim is to assemble in one place a wide range of the available evidence from deontic logic, syntax, semantics, normative ethics, and metaethics that bears on these questions and at least attempt to see the forest, over the particular concerns of the linguists, logicians, and moral philosophers who have been interested in *some* but not *all* of the relevant issues bearing on this question. In contrast to the Chisholm/Williams/Kratzer view, I will be arguing in part 2 that a wide range of evidence of many different kinds supports not only the view that there is a sense of ‘ought’ that expresses a relation, one of whose argument-places is for an agent, but also a particular view of its *syntax*.<sup>5</sup> With respect to this first debate, Broome and Wedgwood are on my side, but because they do not accept this view of the syntax of the deliberative ‘ought’, they are limited to a more restrictive range of evidence. Then in part 3, I will argue that the Broome/Wedgwood view on which the deliberative ‘ought’ relates agents to *propositions* is too flexible. In contrast, I will be arguing that the deliberative ‘ought’ relates agents to *actions*, interpreted as a kind of *property* of agents. Finally, in part 4, I will get to part of the payoff, by explaining a range of issues from metaethics, practical reason, and normative ethics where arguments and theories depend heavily on which view about the argument structure of the relation expressed by ‘ought’ is correct. The moral is that the questions tackled in this essay are not of merely arcane interest but go to the heart of a striking range of questions across different areas of moral philosophy.

5. Though I call this a “particular view of its syntax,” the view in this essay is not specific with respect to many questions of interest to linguists; it merely types evaluative and deliberative senses of ‘ought’, respectively, with two different types of syntactic phenomena, the relationship between which is a further theoretical question in linguistic syntax. What will be important for us in this essay is that the evaluative and descriptive ‘ought’ sentences *pattern differently* with respect to their syntax, not the particular underlying syntactic theory about what makes them pattern differently.

## 1.2. Issues and Concepts

Before we get further, it will help to have some useful distinctions and terminology on board. As acknowledged earlier, in addition to the uses of interest to normative ethics, ‘ought’ also functions in English as an epistemic modal, expressing a concept of epistemic likelihood. According to some views, this is simply the result of the kind of outright ambiguity we see with ‘bank’;<sup>6</sup> according to others, the epistemic and evaluative or deontic uses of ‘ought’ result from a single underlying meaning, in the very robust sense that either they actually have the very same truth-conditions, or differences in their truth-conditions result merely from different settings on some contextual parameter.<sup>7</sup> I will not be taking any view about this question in this essay but will be setting aside the epistemic ‘ought’ for the remainder of the essay. Henceforward, all uses of ‘ought’ under consideration are to be understood as *normative* uses.

Among normative uses of ‘ought’, it should be agreed on all sides that some do not relate agents to actions. Paradigms include such sentences as ‘There ought to be world peace’, ‘The meeting ought to start at noon’, and ‘Things ought to improve’. This sense of ‘ought’ has been referred to as ‘ought to be’ (in contrast to ‘ought to do’), and Wedgwood 2007 follows Sidgwick 1907 in referring to it as the ‘political ought’. I don’t find either of these terms wholly satisfactory; for lack of a better word, I will be dubbing this the *evaluative* sense of ‘ought’, by which I do not mean to import any particular theory of its meaning or truth-conditions.

The controversy between all of the views at issue in this essay arises with respect to sentences that say of some agent and some action, that he or she ought to do it—sentences like ‘Jim ought to jam’, ‘Sal ought to sail’, and ‘Bill ought to bail’. I will call all such sentences *agential ‘ought’ sentences*. Note that here our terminology must be more cautious because much more is controversial: the Chisholm/Williams/Kratzer view holds that agential ‘ought’ sentences *all* express the same evaluative sense of ‘ought’ as other sentences do, and so I will refer to that view as the *semantic*

6. Compare Thomson 2007, though both cross-linguistic evidence and epistemic/evaluative ambiguities for a range of other modals in English make it very difficult to maintain that the ambiguity is a coincidence, as with ‘bank’.

7. Compare recent work by Stephen Finlay (forthcoming, n.d.), who assigns the epistemic and evaluative uses of ‘ought’ the very same truth-conditions and the classic contribution of Kratzer (1977), for whom their truth-conditions differ only because of the value of a contextually set parameter.

*uniformity thesis*. In contrast, the naive view holds that *some* uses of agential ‘ought’ sentences express a distinct relation between agents and actions—a relation that I will call, again for lack of a better term, the *deliberative* sense of ‘ought’. Crucially, however, we will see in the next section that proponents of defensible versions of the naive view must admit that due to the systematic ambiguity of ‘ought’ between evaluative and deliberative senses, agential ‘ought’ sentences are actually systematically ambiguous between deliberative and evaluative readings.

### 1.3. The Evaluative ‘Ought’: Syntax and Semantics

Even setting aside the epistemic ‘ought’, we should all be able to begin by agreeing on the existence of what I have been calling the *evaluative* sense of ‘ought’. The evidence comes from sentences like the following:

- 1a** There ought to be world peace.
- 1b** World peace ought to obtain.
- 1c** The meeting ought to start at noon.
- 1d** It ought to be that the meeting starts at noon.
- 1e** It ought to be that Jim jams.

None of these sentences is plausibly understood as expressing a relation between an agent and an action, but each *is* plausibly interpreted as making a broadly normative claim, in contrast to a merely epistemic claim about what is likely to be the case. So they are examples of evaluative uses of ‘ought’.

Noting that pairs like **1c** and **1d** appear to be trivially equivalent, it is natural to compare them to other pairs, such as the following:

- 2a** The meeting seemed to start at noon.
- 2b** It seemed that the meeting started at noon.
- 2c** The meeting is likely to start at noon.
- 2d** It is likely that the meeting will start at noon.

Each of these pairs of sentences appear to be trivially equivalent, for which linguists have a simple explanation: each member of each pair is associated with a single underlying structure at the level of semantic interpretation, and the members of the pairs differ only in how they satisfy the requirements of English grammaticality. Very roughly, on this view the underlying structure that these sentences share looks like this:

- 2e** [seemed][[the meeting][starts at noon]]

But whereas English sentences require a grammatical subject, this structure does not provide a grammatical subject (it starts with a verb, ‘seemed’). It turns out that there are essentially two ways of getting around this problem, one of which involves “moving” the subject of ‘starts at noon’ up to be pronounced as the subject of ‘seemed’, and one of which is to insert a nonreferential or *expletive* ‘it’ into the subject place. The first members of each pair satisfy the requirements of grammaticality in the former way, and the second members satisfy it in the latter way—but the fact that they have the same (or at least closely related) logical form explains why they are trivially equivalent.

As sentences **1a–1e** illustrate, the evaluative sense of ‘ought’ falls into the same group as ‘seemed’ and ‘is likely’; it is what linguists call a “raising” verb. Raising verbs all share the feature that their subject-places are semantically null and are filled either by a nonreferring ‘it’ or ‘there’, as in the examples above, or by a noun-phrase that “raises” from a lower clause in order to make the sentence grammatical.<sup>8</sup> I’ll introduce a set of tests for raising verbs in section 2.4, but for now it suffices to understand what they are.

Raising verbs are all naturally semantically interpreted as expressing propositional *operators*; such operators may be simple, context-invariant operators applied to the proposition expressed by their *pre-jacent*, modeled on ‘it is necessary that’, or they may require further arguments that are supplied only by context—for example, interpreting ‘The meeting seemed to start at noon’ seems to require knowing *to whom* it seemed that the meeting started at noon. The answer to this question can be supplied only by context. The Kripke semantics for Standard Deontic Logic treats ‘ought’ as a contextually invariant propositional operator, OUGHT(P), making it more like ‘necessarily’; Angelika Kratzer’s semantics requires two contextually supplied argument-places, making it more like ‘seemed’.

As the examples of ‘seemed’ and ‘is likely’ illustrate, it is no consequence of the fact that ‘ought’ is a raising verb or is naturally interpreted

8. It is not important for what follows exactly how the underlying story from linguistic syntax about ‘raising’ verbs is supposed to go. What *is* important in what follows is: (1) raising verbs apply indiscriminately to any sentence—including agential sentences, (2) the subject place of raising verbs is semantically inert, and (3) raising verbs yield equivalences between sentences with ‘it’ or ‘there’ and sentences with ‘raised’ subjects. More fine-grained differences in how these facts are explained will not play a role here.

as expressing a propositional operator that it does not express a relation one of whose argument-places is for an agent. To interpret the sentence ‘the meeting seemed to start at noon’, we must know to whom it seemed that the meeting started at noon. Similarly, it is plausible that the proposition that the meeting will start at noon is not likely or unlikely *simpliciter* but only relative to a set of background information—perhaps relative to some *person’s* background information. So ‘seemed’ *does* express a relation, one of whose places is for an agent—at least, for an experiencing subject—and ‘is likely’ can be plausibly interpreted as doing so, as well. So if ‘ought’ is like ‘seemed’, then it could express a relation between agents and propositions.<sup>9</sup> Nevertheless, evaluating the truth of ‘there ought to be world peace’ does not seem to require knowing to whom or for whom or by whom there ought to be world peace. And so the evaluative ‘ought’ is most naturally interpreted as expressing a propositional operator that does *not*, as ‘seemed’ does, also require an agent for an argument.

It is an important consequence of the raising syntax of the evaluative ‘ought’ that, given general principles about free combination, agential ‘ought’ sentences can also be understood as involving the evaluative ‘ought’, as illustrated by the following sentences, each paired with their ‘seemed’ counterpart:

- 3a** Jim ought to jam.
- 3b** Jim seemed to jam.
- 3c** Sal ought to sail.
- 3d** Sal seemed to sail.
- 3e** Bill ought to bail.
- 3f** Bill seemed to bail.

As these sentences show, raising verbs allow us to generate agential sentences. The very same syntactic mechanism (whatever it is) that allows ‘the meeting’ to raise to the subject-position of ‘the meeting ought to start at noon’ allows ‘Jim’ to raise to the subject-position of ‘Jim ought to

9. In fact, in both Broome’s and Wedgwood’s views, this is how the deliberative sense of ‘ought’ works—it is a raising verb expressing a propositional operator whose interpretation requires a contextually supplied subject. This interpretation is also consistent with standard approaches to deontic logic, so long as the subject is assumed to be held fixed and suppressed in the formalism.



jam’.<sup>10</sup> This is why I was careful to distinguish talk about agential ‘ought’ sentences from talk about *deliberative uses* of agential ‘ought’ sentences, rather than running them together under talk about ‘ought-to-do’, as the literature has been wont to do. Proponents of the naive view cannot claim simply that the meaning of sentences like ‘Jim ought to jam’ cannot be reduced to the evaluative ‘ought’—for as these comparisons with ‘seemed’ illustrate, straightforward principles about free recombination allow the evaluative ‘ought’ to generate sentences like these. So what careful proponents of the naive view must say is that agential ‘ought’ sentences are systematically *ambiguous* between deliberative senses and evaluative senses. That is what I will be arguing in part 2.

### 2.1. Getting Comfortable with the Deliberative ‘Ought’: Five Hallmarks

As we noted in section 1.4, it is a consequence of the raising syntax of the evaluative ‘ought’ that every sentence of the form ‘ $x$  ought to  $A$ ’ that is formed using the evaluative ‘ought’ has an equivalent counterpart of the form ‘it ought to be that  $x$   $As$ ’. As we noted, this is a result of the fact that these two kinds of sentence simply satisfy the requirements of English grammaticality in different ways, but we also confirmed this observation by comparison with ‘seemed’ and ‘is likely’. Consequently, the way to argue for a distinctive *deliberative* sense of ‘ought’ is to look for sentences of the form ‘ $x$  ought to  $A$ ’ that are *not* equivalent to the corresponding sentence ‘it ought to be that  $x$   $As$ ’. In fact, there are a number of reasons to suspect that many agential ‘ought’ sentences do not obey this equivalence.

Consider the case of Luckless Larry, who has recently come by many misfortunes—his parents and siblings have recently passed away, his wife has divorced him to run off with a younger man, he has lost his job including his health insurance, and he has recently been diagnosed with kidney disease, which will require expensive treatment. Larry deserves to win the lottery, if anyone does. So if there is to be any justice in this world, then Larry ought to win the lottery. It ought to be that Larry wins the lottery. These two claims seem to be equivalent—which is what we’ve just seen that we would predict if they both involve the evaluative ‘ought’.

10. Just to reiterate: the important feature of raising verbs in which I am interested is that they generate these equivalences, not the exact mechanism by which this happens. I’ll continue to describe things in the main text in terms of ‘raising’ or ‘movement’, and linguistically sophisticated readers should translate those claims into their preferred framework if necessary.

On the other hand, if Larry comes to you seeking advice about what to do, you are not likely to tell him that he ought to win the lottery. Moreover, the reason why you are not likely to tell him this is not simply that it is not relevant—after all, the very question that Larry comes to you with is the question of what he ought to do. So if ‘Larry ought to win the lottery’ is *unambiguously* an evaluative ‘ought’ sentence, then it is both true and relevant. This suggests that what Larry is interested in, when he comes to you for advice, is *not* the evaluative sense of ‘ought’—but rather something else: the ‘ought’ of advice, or deliberation—the *deliberative* ‘ought’. Two important hallmarks of the deliberative ‘ought’ are that it matters directly for *advice* (compare MacFarlane and Kolodny 2010) and is the right kind of thing to *close deliberation* (compare Ross 2010).<sup>11</sup>

When someone comes to you for advice about what he or she ought to do, one of the relevant conditions on the correct response is whether he or she is accountable if he or she does not do it—it is legitimate criticism of someone that he or she does not do what he or she ought to have done, in the sense of what it was advisable for him or her to do. But Larry is not accountable if he does not win the lottery. Nor is the meeting accountable if it does not start at noon, even though the meeting ought to start at noon. So a third hallmark of the deliberative ‘ought’ is that when someone *deliberatively* ought to do something, he or she is accountable if he or she does not do it, whereas this is not true in general if it is merely true if it evaluatively ought to be that he or she does it. The accountability hallmark of the deliberative ‘ought’ is developed and defended by Broome n.d.

A fourth hallmark of the deliberative ‘ought’ illustrated by the case of Luckless Larry is the principle that ‘ought’ implies ‘can’. When someone comes to you for advice about what he or she ought to do, answers that are beyond his or her ability are also inappropriate—the ‘ought’ of advice and deliberation implies ‘can’. Not just ‘can’ in the sense of bare

11. When I say that the deliberative ‘ought’ matters *directly* for advice and *closes* deliberation, I mean to say more than that it is relevant for advising someone what to do or that it is one factor relevant in deliberation. I mean that knowing what someone ought to do, in the deliberative sense, settles the question of what it is advisable for them to do and that knowing what one ought to do, in the deliberative sense, settles the question of what to do, rather than simply being one important factor among others. This point is very important—arguments in the literature often trade on supposed deliberative readings of evaluative ‘ought’ sentences that are *indirectly* relevant for advice and deliberation. See, for example, Bhatt 1998.

possibility but ‘can’ in the sense of *ability*; it has to actually be in his or her power to *do*. But clearly winning the lottery is not something that Larry is able to do. That in itself seems to be sufficient to rule it out as the answer to what he ought to do that is appropriate when he seeks advice from you, but it clearly doesn’t rule out our initial conclusion that Larry *ought* to win the lottery—because he deserves it. The idea that the principle that ‘ought’ implies ‘can’ applies to the deliberative ‘ought’ but not the evaluative ‘ought’ was discussed by both G. E. Moore (1922) and Lloyd Humberstone (1971) and is appealed to by Ralph Wedgwood (2007).

A final hallmark of the deliberative ‘ought’ can be illustrated by comparison to *obligations*. Contrary to many deontic logicians’ descriptions of their study of the logic of ‘ought’ as the logic of obligation, the notion of obligation is both stronger and weaker than that expressed by the deliberative ‘ought’. It is stronger because it can be the case that someone ought to do something even though he or she is not obligated to do it because obligations are *strict* in a way that ‘ought’ is not. On the other hand, it is weaker because, since obligations can conflict but ‘ought’ is an all-things-considered, you can have an obligation to do something, even though it is not the case that you ought to do it (because you have a weightier contrary obligation). Nevertheless, the deliberative ‘ought’ is *more closely connected* to the notion of obligation than the evaluative ‘ought’. When we say that Larry ought to win the lottery because he deserves to, we are intuitively saying nothing like that Larry has an obligation to win the lottery, whereas when we say in response to Larry’s request for advice that it is not the case that Larry ought to win the lottery—because it’s not under his control—it is intuitive that we are denying anything like an obligation of Larry to win.

More evidence is required for anything like a proof, but the case of Larry illustrates the naive view very well. It shows that there are *some* uses of agential ‘ought’ sentences, like ‘Larry ought to win the lottery’, which is transparently equivalent to ‘It ought to be that Larry wins the lottery’, that are not tied to advice and do not close deliberation, failure to comply with which does not imply accountability of the agent, that do not imply ‘can’, and that are in no way connected to obligation. But they also suggest that there are *other* uses of agential ‘ought’ sentences, like ‘Larry ought to win the lottery’, which is *not* transparently equivalent to ‘It ought to be that Larry wins the lottery’, that *are* tied directly to advice and *do* close deliberation, failure to comply with which *does* imply accountability of the agent, that *do* imply ‘can’, and that intuitively make claims that are similar in kind to claims of obligation. According to the naive

view, these latter, advice- or deliberation-oriented readings of agential ‘ought’ sentences arise from a distinct, *deliberative* sense of ‘ought’, so that agential ‘ought’ sentences like ‘Larry ought to win the lottery’ are systematically ambiguous between evaluative and deliberative readings.<sup>12</sup> The foregoing remarks should illustrate how this idea works and where it comes from. But there is much more to be said in favor of it.

## 2.2. From Ambiguity in Sentences to Ambiguity in ‘Ought’

If we grant that the sentence ‘Larry ought to win the lottery’ is genuinely ambiguous between a reading on which it is true in Larry’s case because he deserves to win the lottery and a reading on which it is false in Larry’s case because he is not able to win the lottery, we are on our way to justifying the thesis that ‘ought’ is ambiguous. But we are not there yet for one prominent response to the ambiguity in this sentence is to claim that it arises not from an ambiguity in ‘ought’ but from an ambiguity in the *rest* of the sentence. This move was suggested by Horty and Belnap 1996 and is defended in Horty 2001.

The idea behind this move is very simple. Given the raising syntax of the evaluative ‘ought’, ‘Larry ought to win the lottery’ is generated by applying ‘ought’ to a *prejacent* sentence, ‘Larry wins the lottery’. So if this prejacent sentence is ambiguous, then ‘Larry ought to win the lottery’ will also be ambiguous—even without postulating any ambiguity in ‘ought’. Horty and Belnap motivate suspecting that ‘Larry wins the lottery’ is at least in principle ambiguous, by observing that *some* NP-VP (noun-phrase-verb-phrase) sentences are *agential*, in the sense that they say that some agent *does* something, whereas others are merely *circumstantial*, in the sense that they say merely that something *happens* to someone. An agential reading of ‘Larry wins the lottery’ would credit winning the lottery as something that Larry did, whereas a circumstantial reading of ‘Larry wins the lottery’ would merely report that it was something that happened to Larry.

The idea that some NP-VP sentences are agential and some are circumstantial is a common one in linguistic syntax and semantics and is associated with the idea that the argument-places of verbs come marked with *thematic roles* or *theta-roles*, which specify whether the subject, for example, is an AGENT or an EXPERIENCER in the relation.<sup>13</sup> Logics of

12. Again, previous caveats about ‘ambiguous’ apply.

13. See, for example, Jackendoff 1972; Grimshaw 1990; and E. Williams 1994.

agency, including particularly that developed in Belnap, Perloff, and Xu 2001, provide a way of interpreting the semantic significance of the difference between *agential* and *circumstantial* sentences by treating circumstantial sentences as basic and introducing an agency operator (abbreviated ‘stit’ for ‘sees to it that’) that transforms circumstantial sentences into corresponding agential sentences. So, for example, linguists would distinguish, at least in principle, between ‘Larry<sub>AGENT</sub> wins the lottery’ and ‘Larry<sub>EXPERIENCER</sub> wins the lottery’, and the semantic significance of this can be captured in a logic for agency by distinguishing between ‘Larry stit: Wins(Larry)’ and ‘Wins(Larry)’.<sup>14</sup>

Moreover, the agency-in-the-prejacent theory can actually predict some of the hallmarks of the deliberative uses of agential ‘ought’ sentences. For example, if we assume that in general ‘it ought to be that P’ implies the possibility of P but not the ability of anyone to bring P about and understand the relevant sense of ‘can’ as ‘it is possible for him or her to bring it about that’, then since on the deliberative reading, ‘Larry ought to win the lottery’ means ‘It ought to be that Larry brings it about that Larry wins the lottery’, this will imply that Larry is able to bring it about that he wins the lottery—so we get the prediction that ‘ought’ implies ‘can’ applies to deliberative ‘ought’ sentences but not in general—without postulating any ambiguity in ‘ought’.<sup>15</sup> All of this is to say that the hypothesis that agential ‘ought’ sentences are ambiguous because their prejacent are ambiguous is initially well motivated, and so it presents a serious challenge to the naive view. For ease of reference, I’ll call it the *agency-in-the-prejacent* theory.

The agency-in-the-prejacent theory does well at distinguishing the two possible readings of ‘Larry ought to win the lottery’ because it is at least initially plausible that ‘Larry wins the lottery’ is, in fact, ambiguous between agential and circumstantial readings. But once we consider other cases, the agency-in-the-prejacent theory both overgenerates and undergenerates; I’ll cover overgeneration first, then undergeneration.

14. Compare Belnap, Perloff, and Xu 2001 and Horty 2001. Note that logicians working on the logic of agency, although inspired by the distinction between agential and circumstantial readings of natural language sentences, have not actually been trying to predict the behavior of natural language sentences, much less to do so in a way that respects how those sentences are actually composed. Nevertheless, even though this has not been the logicians’ primary interest, ultimately a logic of agency tells us something about agency in natural language only if there is some way of mapping it to the composition of natural language sentences.

15. See the discussion in Horty 2001 for details.

*Overgenerating: Agency-in-the-Prejacent*

The agency-in-the-prejacent theory overgenerates because it is a consequence of the agency-in-the-prejacent theory that the *very same* ambiguities that we observe in ‘Larry ought to win the lottery’ can be observed in ‘It ought to be that Larry wins the lottery’. But our five earmarks for deliberative ought’s don’t arise for ‘It ought to be that Larry wins the lottery’. The question of whether it ought to be that Larry wins the lottery simply isn’t relevant for advising Larry about what to do, and it doesn’t settle the deliberative question for Larry about what *to* do. It does not follow from that fact that Larry is unable to win the lottery that it is not the case that it ought to be that he does, and the claim that it ought to be that Larry wins the lottery is not plausibly construed as on a par with the claim that he is obligated to win it. So even if the agency-in-the-prejacent theory is right that ‘Larry wins the lottery’ is ambiguous, it can’t be that ambiguity that generated our observations about Larry in the first place.

Similarly, the agency-in-the-prejacent theory predicts the same ambiguities in ‘Larry seemed to win the lottery’ and ‘It seemed that Larry won the lottery’ as in ‘Larry ought to win the lottery’—because these involve applying ‘seemed’ to the very same prejacent, which if the agency-in-the-prejacent theory is true, is ambiguous in the very same ways. But ‘Larry seemed to win the lottery’ does not appear to be subject to any similar phenomena to that we observed about ‘Larry ought to win the lottery’. Circumstantial evidence for this is that despite the fact that philosophers have perennially found the naive view of ‘ought’ to be intuitive, no one seems to have found any analogous view about ‘seemed’ intuitively compelling. Nor are there anything like the list of hallmarks for the deliberative ‘ought’ that arise for some readings of ‘Larry seemed to win the lottery’ but not for others.

*Undergeneration: The Passivization Test*

The agency-in-the-complement theory also undergenerates because it is possible to observe the very same ambiguity that we observe in ‘Larry ought to win the lottery’ in sentences in which it is not at all plausible to make the required claims about possible readings of the prejacent. The classic—and best—example is an old one (introduced by Gilbert Harman [1973] and discussed by P. T. Geach [1982], who attributes the point to Anselm):

- 4a** Bill ought to kiss Lucy.  
**4b** Lucy ought to be kissed by Bill.

These two sentences intuitively say quite different things—and in this they are much more like ‘wants’ than ‘seemed’:

- 4c** Bill wants to kiss Lucy.  
**4d** Lucy wants to be kissed by Bill.  
**4e** Bill seemed to kiss Lucy.  
**4f** Lucy seemed to be kissed by Bill.

**4e** and **4f** are transparently equivalent, and **4c** and **4d** are transparently nonequivalent. The transparent equivalence of **4e** and **4f** comes from the fact that due to the raising syntax of ‘seemed’, **4e** is equivalent to ‘It seemed that Bill kissed Lucy’, **4f** is equivalent to ‘It seemed that Lucy was kissed by Bill’, and ‘Bill kissed Lucy’ and ‘Lucy was kissed by Bill’ are transparently equivalent. (The transparent nonequivalence of **4c** and **4d** comes from the fact that *neither* is equivalent to ‘It wants that Bill kisses Lucy’ or ‘It wants that Lucy is kissed by Bill’—more on this in the next section.)

Now, to be careful, on the *evaluative* sense of ‘ought’, **4a** and **4b** are equivalent. Suppose, for example, that Bill and Lucy are characters in parallel stories told in the same film, neither of whom knows the other, but each of whom goes through a series of ill-fated romances that, in parallel, demonstrate how well suited they would be for one another (alas that they never meet!). A natural thing to say about the characters in this imagined film is, ‘Bill ought to kiss Lucy!’. An equally natural thing to say is ‘Lucy ought to be kissed by Bill!’. In this scenario, these two sentences are naturally understood as expressing the very same thought—namely, that these two lovable but romantically ill-fated characters ought to get together. We know that this is the evaluative sense of ‘ought’ not only because these readings both seem equivalent to the claim that it ought to be that Bill kisses Lucy but also because neither is accountable, on account of the fact that they do not actually kiss.

Nevertheless, putting aside the reading on which **4a** and **4b** are equivalent, there is also very clearly a reading on which they are *not* equivalent—in stark contrast to **4e** and **4f**. In particular, **4a** has an admissible deliberative reading—on which it is appropriate for advice for Bill, settles Bill’s deliberative question, implies that Bill is accountable if he does not kiss Lucy, implies that Bill is able to kiss Lucy, and is on a par with the claim that Bill has an obligation to kiss Lucy. But **4b** has no equivalent

reading—on which it is appropriate for advice for *Bill*, settles *Bill*'s deliberative question, implies that *Bill* is accountable if he does not kiss Lucy, implies that *Bill* is able to kiss Lucy, or is on a par with the claim that *Bill* has an obligation to kiss Lucy. In other words, one of the important data about this example is that **4a** has a deliberative reading that is *unavailable* for **4b**.

For the agency-in-the-prejacent theory to explain this, it must assume that the prejacent of **4a**, 'Bill kisses Lucy', has a reading that is unavailable for the prejacent of **4b**, 'Lucy is kissed by Bill'. In particular, it must assume that 'Bill kisses Lucy' has an agential reading on which Bill is the agent: 'Bill<sub>AGENT</sub> kisses Lucy', but that 'Lucy is kissed by Bill' has *no* agential reading on which Bill is the agent. But this is completely implausible. 'Bill kisses Lucy' and 'Lucy is kissed by Bill' differ only by the passive transformation, which *preserves* whether Bill was the agent of the kissing.<sup>16</sup> So the assumption that the agency-in-the-prejacent theory requires in order to deal with this case is simply not true. Moreover, even if it was, it would equally distinguish between **4e** and **4f** since they have the same prejacent. This would allow for a reading of 'Bill seemed to kiss Lucy' for which there is no equivalent reading of 'Lucy seemed to be kissed by Bill'—which is absurd.

So even though the agency-in-the-prejacent theory has some initial promise, it can't deal with the full range of the data. Consequently, the ambiguity between readings of agential 'ought' sentences on which they are appropriate for advice and deliberation, and readings on which they are not, *does* motivate postulating an ambiguity in 'ought' between deliberative and evaluative senses.

### 2.3. The Deliberative 'Ought': Syntax and Other Linguistic Tests

The passivization test, even though it relies on intuitions about readings of sentences that some people find to be fairly subtle without the help of hallmarks of deliberative versus evaluative uses, is strong evidence that there is a sense of 'ought' on which it expresses a relation between an agent and something else, and that the noun phrase appearing in the subject-position before the 'ought' fills this agent position. On this view of the syntax and semantics of 'ought', there is no puzzle about why there is

16. What is important here is not whether the "passive transformation" is strictly speaking a *transformation*, in the sense of transformational grammar, but the fact that in 'Lucy is kissed by Bill', Bill is still the agent of the kissing.



no reading of ‘Lucy ought to be kissed by Bill’ on which it is suited for advice for Bill and Bill’s deliberation, implies Bill’s accountability if he does not kiss Lucy, or on which Bill must be able to kiss Lucy. There is no such reading because the person to whom the ‘ought’ claim is relevant for advice, who is accountable if he or she does not act, who must be able to act, and to whose deliberation the ‘ought’ claim is relevant is simply whoever appears in the subject-position because that is a real, semantic argument-position of ‘ought’.

On this theory, this sense of ‘ought’ does not involve a raising verb at all (recall that the subject-positions of raising verbs are semantically null), but rather has a semantic argument-position for an agent. When ‘Bill’ appears before the deliberative ‘ought’, he is the one who is filling a thematic role for the ‘ought’. This deliberative sense of ‘ought’ is therefore of a different syntactic category than the evaluative ‘ought’ distinguished in section 1.3.<sup>17</sup> It is what linguists call a *control* verb. Unsurprisingly, the test that is employed by the passivization argument is actually one of the *standard* tests employed by linguists to distinguish raising verbs from control verbs. As we’ve already seen, ‘seemed’ is a raising verb, whereas ‘wants’ is a control verb:

- 4c** Bill wants to kiss Lucy.
- 4d** Lucy wants to be kissed by Bill.
- 4e** Bill seemed to kiss Lucy.
- 4f** Lucy seemed to be kissed by Bill.

To distinguish a raising verb from a control verb, compare sentence pairs like **4c** and **4d**, or **4e** and **4f**. If these sentences are equivalent, then the verb passes the passivization test and is a raising verb—as can easily be verified with ‘seemed’ in **4e** and **4f**. If the sentences are nonequivalent, then the verb fails the test and is a control verb—as can easily be verified with ‘wants’ in **4c** and **4d**. In section 2.2 I tried to explain *why* this test works, but the more general point is that it is in fact a general test used more widely in linguistic syntax.

‘Ought’ is more complicated than either ‘wants’ or ‘seemed’, but as I have been arguing, that is only because it is ambiguous between

17. Again, what is important for our purposes here is simply that the evaluative ‘ought’ and the deliberative ‘ought’ *pattern* differently, in ways that a deeper syntactic theory will go on to explain. This has the concrete consequence that the ‘agent’ of a deliberative ‘ought’ sentence needs to appear in the subject-place, rather than simply needing to be contextually salient, which is all that is important for the following tests.

a raising verb and a control verb. Sentences that employ the evaluative ‘ought’ pass the passivization test because they have raising syntax, but sentences that employ the control ‘ought’ fail it—as we’ve seen—because they have control syntax. We require no special apparatus in order to explain this; it arises simply because ‘ought’ has both a sense that behaves like ‘seemed’ and a sense that behaves like ‘wants’. When I say that ‘ought’ is ‘ambiguous’ between these two ‘senses’, I just mean this: even when we exclude ‘epistemic’ readings, the word ‘ought’ sometimes exhibits raising syntax and sometimes exhibits control syntax, and its semantic significance is different in each of these cases—each of the hallmarks of the deliberative ‘ought’ is exhibited in all and only the cases exhibiting control syntax.

This dual patterning of the evaluative versus the deliberative ‘ought’ is important to appreciate. It is supported by the remainder of a battery of tests used by linguists to distinguish raising verbs from control verbs. For example, (1) only raising verbs admit of expletive subjects like ‘it’ or ‘there’, (2) only raising verbs admit idiomatic subjects, and (3) control verbs place restrictions on what sorts of subjects can be allowed without anomaly.<sup>18</sup> Take the expletive subject test first:

- 5a** It seemed/\*wanted/ought to be assumed that he is capable.  
**5b** There seemed/\*wanted/ought to be world peace.

As **5a** and **5b** illustrate, ‘seemed’ works fine in a sentence whose subject is the expletive (nonreferential) ‘it’ or ‘there’, but ‘wanted’ doesn’t—this is because the subject-position of ‘wanted’ is a real, semantic argument-position, and so ‘wanted’ sentences make sense only if their subject-argument has a referent to fill the *agent* place of the WANTED relation. This test predicts that ‘ought’ sentences with expletive subjects will be acceptable but will only admit of evaluative readings—a prediction that I believe is correct: ‘It ought to be assumed that he is capable’ and ‘There ought to be world peace’ do not plausibly imply ‘can’ and do not imply anyone’s accountability; they are not directly linked to advice and do not close deliberation.

Next, take the idiomatic subject test:

- 5c** All hell seemed/\*wanted/ought to break loose.  
**5d** The cat seemed/\*wanted/ought to get his tongue.

18. I borrow these tests from Radford 2004, 268–74.

As **5c** and **5d** illustrate, ‘seemed’ can ‘split’ a noun-verb phrase idiom like ‘all hell broke loose’ or ‘the cat got his tongue’ and still generate an idiomatic reading, but if ‘wanted’ is inserted in the middle of such an idiom, it admits only of a literal, nonidiomatic reading. This is because the proper parts of an idiom have no semantic significance in their own right, but control verbs like ‘wanted’ require a subject-argument that *does* have semantic significance. This test predicts that sentences in which ‘ought’ splits such an idiom will be acceptable but will admit only of evaluative readings—again, I think correctly: ‘The cat ought to get his tongue’ cannot be read as both idiomatic and as expressing an intuitive deliberative ‘ought’ claim, on which it is appropriate for giving the cat advice, the cat is accountable if it does not get his tongue, and so forth. Any such reading requires taking ‘the cat’ to actually refer to a cat, which is inconsistent with the idiomatic reading.

Finally, take the thematic role test:

- 5e** The meeting seemed/\*wanted/ought to start at noon.
- 5f** Jerusalem seemed/\*wanted/ought to be divided between Israel and Palestine.
- 5g** Fermat’s Last Theorem seemed/\*wanted/ought to be provable.
- 5h** Yesterday seemed/\*wanted/ought to have been forgotten.

As **5e–5h** illustrate, ‘seemed’ allows for a subject of any kind—meetings, cities, theorems, and days are all admissible, for example. But in contrast, sentences involving ‘wanted’ don’t make sense unless their subject is the kind of thing to have psychological states—which meetings, cities, theorems, and days are not. Again, this is because the subject-position of control verbs is semantically significant, but the subject-position of raising verbs is semantically null. Consequently, the subject of ‘wanted’ needs to be of the right category to match the thematic role of the subject-position and the category of the appropriate place of the WANTED relation, but no such requirement is in place for ‘seemed’. This test predicts that ‘ought’ sentences with arbitrary subjects will be admissible, but only ‘ought’ sentences whose subjects are *agents* or are appropriate recipients of advice will admit of control, *deliberative* readings. Again, I think this prediction is correct: meetings, cities, theorems, and days are not accountable, they do not have to be able to do things, and so on.

What I have been arguing here is that not only can we draw the conclusion that ‘ought’ is ambiguous between evaluative and deliberative senses, but we can and should actually draw the conclusion that these

two senses of ‘ought’ pattern differently with respect to their *syntax*—merely overlapping in the phonetic or apparent surface form of sentences like ‘Larry ought to win the lottery’ or ‘Bill ought to kiss Lucy’. Not only does this allow us to make sense of the different ways in which we use agential ‘ought’ sentences in situations of advice and deliberation in comparison to situations of evaluation of outcomes, it is consistent with the key tests used to distinguish raising and control verbs. The only complication with ‘ought’, compared with other verbs like ‘seemed’ and ‘wants’, is that it has *both* a sense as a raising verb (the evaluative ‘ought’) and a sense as a control verb (the deliberative ‘ought’).

#### 2.4. The Deliberative Ought: Logic

The foregoing arguments for the existence of a distinct, deliberative sense of ‘ought’ are supported by a range of distinct observations about the inferential relations between apparent deliberative ‘ought’ claims. Many of these observations about the inferential relations between apparent deliberative ‘ought’ claims are hard to reconcile with the combination of the hypothesis that they really are evaluative ‘ought’ claims and plausible assumptions about the logic of the evaluative ‘ought’.

For example, as Broome (n.d.) observes, different agents can have conflicting responsibilities. He imagines a case in which it is Father Murphy’s job to baptize everyone in the parish who needs to be baptized, Colleen is in the parish and needs to be baptized, but it is in Colleen’s interests to be baptized by the holiest priest she can find—who is Father O’Grady, not Father Murphy. According to Broome, it is plausible that Father Murphy ought to baptize Colleen, but also plausible that Colleen ought to see to it that she is baptized by Father O’Grady (and hence not by Father Murphy). If both of these claims are analyzed in terms of the evaluative ‘ought’, however, then the former says that it ought to be that Father Murphy baptizes Colleen and the latter says that it ought to be that Colleen sees to it that she is not baptized by Father Murphy. But it is *prima facie* implausible that inconsistent things both ought to be the case, and it can’t be both that Father Murphy baptizes Colleen and that Colleen sees to it that she is not baptized by Father Murphy.

Now admittedly, some philosophers have been driven to conclude that it *is* possible for inconsistent things to both be such that they ought to be the case—on the basis of the assumption that individuals sometimes ought to do inconsistent things. For example, some philosophers believe that Sartre’s young Frenchman both ought to join the Free French (for

patriotic reasons) and ought not to join the Free French (in order to stay home with his mother). It is true that this can lead to the conclusion that inconsistent things both ought to be the case *but only on the assumption* of the semantic uniformity thesis. If the claim that Sartre's young Frenchman ought to join the Free French is not equivalent to the claim that it ought to be that he joins the Free French, then *intrapersonal* deontic conflicts don't motivate the conclusion that there can be conflicts in the evaluative 'ought' any more than Broome's *interpersonal* conflicts do.

The cases of intrapersonal and interpersonal deontic conflicts are just one kind of example in which maintaining the semantic uniformity thesis forces us to make implausible claims about the logic of the evaluative 'ought'—in this case, forcing us to allow for conflicts of the evaluative ought, such that 'It ought to be that P' and 'It ought to be that  $\sim$  P' are both true. But there are other important cases. For example, standard deontic logic validates the principle of *inheritance* for 'ought'—according to which if B is a necessary consequence of A, and it ought to be that A, then it ought to be that B.<sup>19</sup> But although the principle of inheritance is controversial, it is *more* controversial and controversial in more ways for deliberative 'ought' claims than it is for evaluative 'ought' claims. For example, suppose that it ought to be that Strategic Bomber drops a bomb that decimates the ammunition factory because that will lead to the speediest possible resolution of the war and the fewest casualties in the end. And suppose that this is true even though if Strategic Bomber drops such a bomb, he will unavoidably also be dropping a bomb that will decimate the elementary school. It follows from inheritance that it ought to be that Strategic Bomber drops a bomb that will decimate the elementary school.

In this case, I think, this conclusion is plausible—*if* it really ought to be that he drops a bomb that will decimate the ammunition factory even though that bomb will also decimate the elementary school, then it ought to be that he drops a bomb that will decimate the elementary school. But I don't think that it is plausible that it follows from the fact that Strategic Bomber ought, in the deliberative sense, to drop a bomb that will decimate the ammunition factory that he ought, in the deliberative sense, to drop a bomb that will decimate the elementary school. On the contrary, since dropping a bomb that will decimate the elementary school is only a predictable side-effect of dropping a bomb that will

19. See Hilpinen and Føllesdal 1971 and McNamara 2006 for discussion of standard deontic logic.

decimate the ammunition factory, and not a necessary *means* to it, you should not *advise* Strategic Bomber to decimate the elementary school, and in no sense should finding out that he ought to decimate the ammunition factory settle the deliberative question for him of whether to decimate the elementary school. Strategic Bomber should accept the decimation of the elementary school as a consequence of his actions, but he shouldn't reason toward an intention to do it in his deliberations.<sup>20</sup> The general idea behind this example, of course, is that the deliberative 'ought' may transmit to *necessary means* but not to all necessary consequences, whereas it is at least much more plausible that the evaluative 'ought' transmits to necessary consequences—even those that are not themselves means.<sup>21</sup>

It is possible to assemble quite a list of other plausible claims about deliberative 'ought's' that are jointly inconsistent with the semantic uniformity thesis and with some fairly plausible claims about the logic of the evaluative 'ought'—for further discussion see especially Krogh and Herrestad 1996; Horty 2001; and Ross 2010. The main point that I wish to make here is that the differences between the actual inferential relationships between deliberative 'ought' sentences and the inferential relations that you would expect them to have if the semantic uniformity thesis were true should not be surprising if, as I have been arguing, the semantic uniformity thesis is in fact false.<sup>22</sup>

## 2.5. Williams's Scope Ambiguity Argument

We've now seen a wide range of evidence that supports the naive view that there are both deliberative and evaluative senses of 'ought' over the semantic uniformity thesis, according to which the evaluative sense is the only sense. Before moving on, I should note that we are now in a position to evaluate perhaps the most significant argument for the semantic

20. Compare Millsap n.d.

21. Of course, this is not to take a stand on whether the evaluative 'ought' transmits to *all* necessary consequences; Good Samaritan cases and Ross's paradox present challenges to the full generality of inheritance that are orthogonal to the point that I am making here. See the appendix to Ross 2010 for further discussion.

22. Note that not everyone who distinguishes between the evaluative and deliberative 'ought's allows that they may obey different logical principles; for example, Wedgwood (2007) explicitly endorses the inferential principles of standard deontic logic for both.

uniformity thesis, which is due to Bernard Williams (1981).<sup>23</sup> Williams’s argument begins by observing that sentence **6a** appears to be ambiguous between readings roughly paraphrasable as **6b** and **6c**:

- 6a** Someone ought to tell the boss.
- 6b** It ought to be that someone tells the boss.
- 6c** Someone is such that he or she ought to tell the boss.

Williams’s argument then requires three assumptions: (1) on the reading paraphrasable as **6b**, **6a** employs the evaluative ‘ought’; (2) on the reading paraphrasable as **6c**, **6a** is a deliberative use of an agential ‘ought’ sentence; and (3) the ambiguity in **6a** is of a familiar kind; it is merely the kind of scope ambiguity with which we are all familiar.

The argument then works like this: if (1), the reading of **6a** on which it can be paraphrased as **6b** involves the evaluative ‘ought’, then it can be stated more formally as **6d**, below. And so if (3), the ambiguity in **6a** is due to a mere scope ambiguity, then it must be that the reading on which it can be paraphrased by **6c** can be stated more formally as **6e**, below.

- 6d** OUGHT<sub>evaluative</sub>( $\exists x: x$  tells the boss)
- 6e**  $\exists x: \text{OUGHT}_{\text{evaluative}}(x$  tells the boss)

But then by the assumption (2) that the reading of **6a** on which it can be paraphrased by **6c** is a deliberative use, it follows that I am wrong, and deliberative uses *are* evaluative ‘ought’ claims. **6e** does not, after all, employ a relational sense of ‘ought’; since its difference from **6d** is the result of a mere scope ambiguity, it employs the *same*, evaluative sense of ‘ought’. Williams’s argument is thus designed to show that in at least some cases, the deliberative uses *have* to be understood as evaluative ‘ought’ claims.

This argument is, I think, clever, but totally unconvincing. Though it is clear that **6a** is ambiguous, it is only a theory that tells us that the ambiguity that we observe is merely a scope ambiguity. On the view that I have been defending, **6a** is in fact *three* ways ambiguous. One reading is

23. Probably the other most important argument for the semantic uniformity thesis is based on the motivation to avoid postulating ambiguities—with some support from cross-linguistic coincidence of the words for the evaluative ‘ought’ and the deliberative ‘ought’. Some intriguing reason to doubt this motivation is provided by Nordlinger and Traugott (1997), who distinguish between the control ‘ought’ and the raising ‘ought’ and argue that the control ‘ought’ preceded the raising ‘ought’ in English by several centuries, based on the textual record. See also Romero 2005.

paraphrasable as **6b**, and *two* readings are paraphrasable as **6c**, depending on whether we interpret the ‘ought’ in **6c** as the evaluative, raising ‘ought’ or the deliberative, control ‘ought’. Further, just as *any* sentence of the form, ‘Jack ought to do *A*’, is ambiguous between raising and control readings, but the control readings are more natural, the control reading of **6c** is, I think, more natural than its raising reading.<sup>24</sup>

This means that of Williams’s three assumptions, (1) is true, but both (2) and (3) turn on the ambiguity in **6c**. On the deliberative reading of **6c**, Williams’s assumption (2) is true: that is, after all, the reading for which I proposed to postulate the control ‘ought’! But on the evaluative reading of **6c**, Williams’s assumption is false: I have insisted all along that agential ‘ought’ sentences are ambiguous between evaluative and deliberative readings. Meanwhile, on the evaluative reading of **6c**, Williams’s assumption (3) is true—the ambiguity between the reading of **6a** that can be paraphrased by **6b** and the reading that can be paraphrased by the raising sense of **6c** really is a mere scope ambiguity. But on the deliberative reading of **6c**, this is of course false.

I conclude that Williams’s argument doesn’t tell us anything that we didn’t already know; since there are both evaluative and deliberative senses of ‘ought’, sentences like ‘She ought to tell the boss’ are ambiguous. This yields the prediction that sentences like **6a** will be, in fact, three ways ambiguous, a prediction which is, I think, on reflection correct. I conclude that ‘ought’ really does, in English, on at least some uses, express a relation between agents and something else. The question that leaves us is: what is the something else?

### 3.1. Relating Agents to Actions

John Broome (n.d.) and Ralph Wedgwood (2006) share my view that when we make ‘ought’ claims, we are concerned, in at least central cases, with a relation that holds between an agent and something else. They do not share my view about the syntax of ‘ought’, however, which means that they cannot avail themselves of what I have taken to be the

24. This is important. In proposing the view that agential ‘ought’ sentences are strictly speaking ambiguous between evaluative, raising readings and deliberative, control readings, I am not committed to the view that these two readings are equally eligible. On the contrary, since there is an alternative, equivalent way to make the evaluative ‘ought’ claim—by using ‘it ought to be that’—standard principles of interpretation predict that the deliberative reading will be the most eligible, except in special contexts.



strongest arguments for this view—both Broome and Wedgwood hold that the deliberative sense of ‘ought’, like the evaluative sense, is a raising verb.<sup>25</sup> And they disagree with me about what agents are related *to* by this ‘ought’. Broome and Wedgwood both claim that it relates agents to propositions, but as a proponent of the naive view, this is something I deny. I claim, and in this part of the essay will proceed to argue, that the deliberative ‘ought’ relates agents to *actions* rather than propositions.

The sense of ‘action’ on which I claim that the deliberative ‘ought’ relates agents to actions is very broad. It can be the case that Max ought to believe that *p*, or that Max ought to be saddened by recent events, but believing that *p* and being saddened by recent events are not commonly thought of as actions. Still, for my purposes I will use ‘action’ in a very broad way to refer to the kind of thing that agents can *do*, in a very broad sense. On my view, we can think of actions as a subclass of nontrivial *properties* of agents—properties over which agents have a certain amount of a certain kind of control. If you find yourself uncomfortable with the idea that actions are a kind of property of agents, then you can characterize my view as holding that the deliberative ‘ought’ relates agents to properties rather than actions, and you may take my view to be a greater departure from the naive view than originally advertised. In any case, what my arguments will really show is that the deliberative ‘ought’ relates agents to a kind of property rather than to arbitrary propositions. But I will continue to use the word ‘action’.

My basic argument that the deliberative ‘ought’ relates agents to actions rather than propositions is very simple. The Basic Problem with the propositional view is that it is too powerful. Whenever an agent ought to do something, Broome and Wedgwood will say that he or she stands in the deliberative OUGHT relation to the proposition that he or she does that thing. In this way, their view will capture everything that I think we can and do talk about when we are interested in this relation and everything that I think it makes sense to talk about. But if OUGHT is just a relation that you can stand in to some proposition—for example, the proposition that you exercise daily—then it is a relation of which it makes sense to ask whether you stand in it to arbitrary other propositions—for example, to the proposition that *I* exercise daily. But I don’t think that it makes sense to ask whether you stand in the OUGHT relation to the proposition that I exercise daily. It’s not just that I think it

25. See Wedgwood 2007; Broome n.d.

is *false* that you ought for me to exercise daily; I don't think this question makes any sense; it involves a category mistake. So if it is false, it is false because it doesn't make sense. The view that the OUGHT relation relates agents to propositions is too powerful because it predicts that some things should make sense that don't—it licenses a category mistake. That is why it is wrong. That is the Basic Problem.

Broome and Wedgwood, however, have some arguments for their view and a counterargument to the Basic Problem. In the following sections, I'll exhibit the evidence and argue that except for the Basic Problem, it is inconclusive.

### 3.2. The Scope Ambiguity Argument, Again: Wedgwood

Wedgwood offers an argument for his view that is based on Bernard Williams's scope ambiguity argument for the semantic uniformity thesis. Recall that Williams's argument was based on the observation that **6a** is ambiguous between readings that can be paraphrased as **6b** and **6c**:

- 6a** Someone ought to tell the boss.
- 6b** It ought to be that someone tells the boss.
- 6c** Someone is such that he or she ought to tell the boss.

Here is what Wedgwood says about a sentence like **6a** (his example is 'go and inform the manager', rather than 'tell the boss'):

Even if one keeps constant the interpretation of 'ought' as having its practical or deliberative sense here, this sentence is clearly ambiguous. The ambiguity is most naturally interpreted as involving a scope ambiguity: on one reading, [6a] means 'It ought to be that: someone goes and informs the manager'; on the other reading, it means 'Someone is such that: *he* ought to go and inform the manager'. On the first reading, the only agent who could possibly be the "subject" of the 'ought' is presumably the group involved in the joint deliberation, viewed as a collective agent. But this collective agent is not explicitly mentioned in the sentence, and so... 'ought' in this first reading of [6a] also seems to be a propositional operator; and as Williams says (1981: 116), "it is hard to see what requires it, or even allows it, to turn into something else" in the second reading. So there seems to be a reason for treating even the practical or deliberative 'ought' as a propositional operator. (Wedgwood 2006, 133–34)

It appears that Wedgwood means to parallel Williams's argument. Like Williams, he assumes (1) that on the first reading, the 'ought'

expresses a propositional operator. Of course, Williams held that it was the evaluative ‘ought’, whereas Wedgwood holds that it is a deliberative ‘ought’, but each starts with the assumption that on the first reading, ‘ought’ takes a propositional argument. Next, like Williams, Wedgwood assumes (2) that the second reading is a deliberative reading. He doesn’t state this explicitly, but it is clear from his reasoning. And finally, like Williams, Wedgwood assumes that the ambiguity in **6a** is due to (“involves”, he says, just to be careful) a scope ambiguity.

However, Wedgwood’s argument differs from Williams’s in more than one way. Most importantly, Wedgwood is not trying to argue for the traditional view that there is only one, evaluative sense of ‘ought’ but for the view that on its deliberative sense, ‘ought’ relates agents to propositions. So if he is to maintain premises (2) and (3), his premise (1) cannot be the same as Williams’s; whereas Williams took the highly plausible view that **6b** involves the evaluative ‘ought’, Wedgwood is forced to hold that **6b** involves a deliberative ‘ought’ but with the whole *group* as the agent. But unfortunately, it is not particularly plausible that **6b** expresses a deliberative ‘ought’ at all. On the contrary, Williams’s original argument got off to a good start only because it seemed uncontroversial that on its first reading, **6a** could *only* involve the evaluative ‘ought’. But this is precisely what Wedgwood needs to deny, in order for *his* argument to get started!

A further problem with Wedgwood’s argument is that given his own view, the ambiguity in **6a** is not, after all, a mere scope ambiguity. For on his view, **6c** can be formalized this way:

- 6f** OUGHT<sub>the group</sub>( $\exists x:x$  tells the boss)<sup>26</sup>  
**6g**  $\exists x$ :OUGHT<sub>*x*</sub>(*x* tells the boss)

The difference between **6f** and **6g** is not merely one of scope; it also turns on how the agent-argument of the ‘ought’ relation is supplied. In **6f** this argument is filled by the group, whereas in **6g** it is filled by no one in particular but is rather bound by the quantifier. Perhaps this is what Wedgwood meant to qualify, when he said that the ambiguity “involves”

26. Notice that since Wedgwood holds that **6b** expresses the deliberative sense of ‘ought’, he is committed to being able to get deliberative readings that the expletive subject test (from section 2.3) classifies as *raising* verbs. This reflects the fact that Wedgwood assumes that the deliberative ‘ought’ exhibits the same syntax as the raising ‘ought’ but simply has an extra argument-position that is supplied *by context*, rather than by an overt syntactic subject argument-position. And this means, as noted earlier, that Wedgwood cannot avail himself of the passivization argument.

a scope ambiguity, but this substantially undercuts the force of his argument. If the ambiguity is not a *mere* scope ambiguity, then that opens the door to other accounts of the nature of the ambiguity.

Moreover, this reveals what is at the heart of the differences between Williams's account, my own, and Wedgwood's over the ambiguity in **6a**. According to Williams, there are two possible readings of **6a**, which differ in terms of the scope of the evaluative 'ought'. According to me, there are both of those readings, plus a further reading, which is given by **6c**, understood as involving the deliberative 'ought'. Whereas according to Wedgwood's view, **6a** is in fact *four* ways ambiguous. Since he allows for the ordinary evaluative 'ought' (which he follows Sidgwick in calling the 'political ought'), he gets both of Williams's readings. And then he gets two more readings, corresponding to the two possible scopes for the deliberative 'ought' and corresponding to whether the group is the agent, or whether the agent-argument is bound by the quantifier.

In fact, since according to Wedgwood the agent-argument of the 'ought' relation is not supplied by any overt argument-place but rather is somehow supplied by the context, as we need in order to generate **6f** as a reading of **6b**, Wedgwood in fact is committed to allowing that **6a** could be used to express as many distinct propositions as there are agents who we might be understood as occupying the agent-argument of the 'ought'. Wedgwood's view is flexible enough to make all of the right distinctions, but it is *too* flexible and generates readings that are not exhibited.

### 3.3. The Argument from Infinitival *To*: Broome

In "Normative Requirements," Broome (1999) took the fact that 'ought' takes an infinitival clause for its complement to justify treating the argument as a proposition.<sup>27</sup> In general, infinitival clauses *do* seem to provide proposition-like arguments. Compare, for example, 'Jon wants to get rich' to 'Jon wants Mary to get rich'. If we interpret 'Jon wants to get rich' as having an unpronounced pronoun referring to Jon, then we can explain the meaning of these two sentences by appeal to the same semantic values and compositional principles. This assumption is

27. To be fair, Broome (1999) may not have understood his remarks about control verbs like 'wants' as an *argument* that 'ought' relates agents to propositions, so much as justifying his *stipulation* that he was going to treat 'ought' in that way. This does not affect the relevance of this sort of argument, which I take to be the most interesting argument in favor of his view.

standard in linguistics, where the postulated unpronounced pronoun in ‘Jon wants to get rich’ is called “PRO.” The reason ‘wants’ is called a “control” verb is that it is said to allow its subject to “control” the PRO argument—which is more or less like binding ‘PRO’ like a variable, so that it essentially refers to the subject of ‘wants’.

This relates to the hypothesis that ‘wants’ relates a being with psychological states to a proposition because when Jon wants Mary to get rich, what Jon wants is for Mary to get rich. He wants the following to obtain: that Mary gets rich. Similarly, when Jon wants to get rich, what Jon wants is for Jon to get rich. He wants the following to obtain: that Jon gets rich. These readings are predicted if the role of the infinitive clause, ‘\_\_ to get rich’, is to semantically contribute a proposition. ‘Mary to get rich’ contributes the proposition that Mary gets rich, ‘PRO to get rich’ contributes the proposition that  $x$  gets rich, where  $x$  is the value of ‘PRO’ determined by the larger sentence in which it figures, and ‘wants’ can be understood as semantically unambiguous between these constructions and constructions like ‘wants that’ ‘wants for’, and so on. So if the role of infinitive clauses is to semantically contribute propositions, and ‘ought’ takes an infinitive complement, that would seem to suggest that ‘ought’ relates agents to propositions—just as ‘wants’ does.

This is an intriguing argument, but unfortunately I believe it goes too fast. Unlike ‘wants’, ‘ought’ allows *only* for a PRO argument. Sentences like ‘Jon ought to get rich’ are fine, but sentences like ‘Jon ought Mary to get rich’ are unacceptable. ‘Ought’, it seems, *selects* for a PRO subject in its complement. Why should this be? I think it is because even though infinitive clauses, in general, have the expressive power to be able to pick out arbitrary atomic propositions, this expressive power is not needed in order to pick out the argument of ‘ought’. The things that Jon ought to do are *actions* and that is why in its control sense, ‘Jon ought’ can be followed only by ‘PRO to  $\varphi$ ’, where ‘ $\varphi$ ’ is a verb phrase denoting some action.

Moreover, significantly, ‘ought’ is not the only verb that takes infinitive clauses but that doesn’t make sense for arbitrary propositional arguments.<sup>28</sup> For example:

- 7a** Maria neglected to show up for the talk.
- 7b** The speaker forgot to bring his notes.
- 7c** Jake proceeded to criticize the speaker’s argument.

28. Special thanks here to Jake Ross and Paul Pietroski for discussion.

- 7d** The speaker will remember to bring his notes next time.  
**7e** Maria aspires to be there next time.

Although Maria can neglect to show up for the talk, she can't neglect Mark to show up for the talk,<sup>29</sup> or neglect for Mark to show up for the talk, or neglect that Mark shows up for the talk. Similarly, although the speaker can forget to bring his notes, he can't forget Maria to bring her notes, or forget for Maria to bring her notes. He can forget *that* Maria brought her notes, of course, but that means something quite different. Again, although Jake can proceed to criticize the argument, he can't proceed Maria to criticize the argument. In each of these cases, the problem isn't just grammar—it's that neglecting is not in general a relation between agents and arbitrary propositions; it's a relation between agents and things that agents *do*—actions. Hence, there is a whole family of control verbs to which the deliberative 'ought' belongs, each of which takes an infinitive clause but does not take arbitrary propositions for its argument. Consequently, Broome's reasoning no more suggests that it makes sense for one person to *ought* for someone else to do something than it shows that it makes sense for one person to proceed for someone else to do something.<sup>30</sup>

29. Note that 'Maria neglected Mark to show up for the talk' is admissible on the reading where it means, 'Maria neglected Mark, in order to show up for the talk', as if neglecting Mark is a means to allowing Maria to attend the talk. This isn't, however, the reading that we need, on which what Maria neglects is the proposition that Mark shows up for the talk.

30. The discussion in this section leaves unanswered exactly how I take the meaning of 'Jim ought to jam' to be composed if 'PRO to jam' in general semantically contributes a proposition, but 'ought' requires a property for its argument. In this note I introduce two possibilities without evaluating them; the same range of options are available for 'neglects', 'forgets', and the others.

The first possibility consonant with my view is that the underlying OUGHT relation is a relation between an agent and an action, but that from this relation we can define a relation between an agent and a proposition, OUGHT<sub>prop</sub>, so that OUGHT<sub>prop</sub>(X,P) just in case for some action A, P is the proposition that X does A, and OUGHT(X,A). Then we say that our language exploits a construction (infinitive clauses) that contributes propositions, in order to allow us to talk about a relation that takes property arguments (the deliberative OUGHT relation), by making OUGHT<sub>prop</sub> the semantic value of the deliberative 'ought'. On this view, as far as semantic compositional principles go, the deliberative 'ought' *does* take a propositional argument, but since the underlying relation is to an action rather than a proposition, we can still explain the category mistakes that the Broome/Wedgwood view allows for. This version of the view is concessive to Broome and Wedgwood at the level of semantics, but nonconcessive at the level of the underlying metaphysics.

### 3.4. The Ungrammaticality Defense

In more recent work, Broome offers an explanation of why ‘Jon ought Mary to get rich’ is ungrammatical. His explanation is that ‘ought’ is not really a control verb like ‘wants’ at all. Broome sets out his explanation in the framework of traditional grammar,<sup>31</sup> but in the syntactic terminology of this essay, his claim is that ‘ought’ does not ever actually have control syntax (never mind his earlier argument just discussed); it has only raising syntax. Unfortunately, because this leaves our language with no way of picking out the relational sense of ‘ought’ that is, as Broome believes, the ‘central normative concept’, we sometimes engage in an activity which he calls “rearsing.” Essentially, “rearsing” involves *pretending* that ‘ought’ is a control verb like ‘wants’, even though it is really a raising verb like ‘seems’.

Since ‘seems’ really is a raising verb and really does have a further, contextually supplied argument place for an agent—just like Broome assumes about the deliberative ‘ought’—it gives us a useful point of comparison for Broome’s hypothesis: what Broome calls “rearsing” would be what happens if we assumed just because ‘Jim seems to jam’ is true or false only relative to *someone* to whom it seems that Jim jams that Jim must be the person to whom this seems to be the case. It is this strange hypothesis that Broome takes to explain why ‘Jon ought Mary to get rich’ is ungrammatical and that leads him to feel justified in stipulatively

---

A second possibility consonant with my view in this essay, inspired by David Lewis’s treatment of *de se* beliefs in Lewis 1979 and by an idea about their relationship to PRO originally due to Chierchia (1989), is to propose that *all* infinitive clauses contribute properties of agents rather than propositions but to identify propositions with ‘pure Cambridge’ properties of agents (of the form  $\lambda x(P)$ , where  $x$  does not appear free in  $P$ ). Then one could add to this view the thesis that PRO obligatorily contributes an agent variable—or in other words that PRO is always bound by the  $\lambda$ -operator. This thesis is consonant with the hypothesis that PRO obligatorily results in *de se* readings, as in ‘Jim wants to jam’. In this framework, my view would be that the deliberative ‘ought’ takes only properties  $\lambda x(P)$ , where  $x$  is free in  $P$ —which is just the plausible claim that ‘being such that  $P$ ’ is not an action—whereas the Broome/Wedgwood view amounts to the thesis that ‘ought’ may take properties  $\lambda x(P)$ , where  $x$  is *not* free in  $P$ . This is of course a very controversial view about the semantic contribution of infinitive clauses, but it actually gives a very elegant compositional picture that suits my view about the deliberative ‘ought’ and intuitively fits very well with ‘neglects’, ‘forgets’, and the others, as well.

31. Broome (n.d., 35–36) takes ‘ought’ to be an auxiliary and understands auxiliaries as combining with main verbs to form a complex verb. This traditional view conflicts with standard tests for constituency—to see which words together form a “unit” in the sentence. See, for example, Radford 2004.

introducing the horrendous ‘Jon ought that Mary gets rich’, in order to be able to express all of the incoherent (as I think) things that he wants to be able to say but are ungrammatical in English. This is the main object of a substantial appendix in Broome n.d.

I say that Broome’s hypothesis is strange because, as I showed in part 2, the evidence is that there *really is* a control sense of ‘ought’—not just raising uses of ‘ought’ that we ‘inexcusably’ “reparse” or pretend have a different syntactic structure and meaning than they really, in fact, do. It is also strange because other cases—including the one that Broome himself discusses at the beginning of Broome 1999—show that no such complicated explanations are needed in order to explain the ungrammaticality of ‘Jon ought Mary to get rich’. For example, ‘hopes’ is, like ‘wants’, a control verb and is so classified by the passive transformation, expletive subject, idiomatic subject, and thematic role tests of section 2.3. Moreover, it also fails Broome’s tests to be an auxiliary verb—the feature that Broome claims creates the grammatical complications with ‘ought’. Yet though ‘Jon hopes to get rich’ is perfectly grammatical, ‘Jon hopes Mary to get rich’ is not. The problem is not that ‘hopes’ belongs to a completely different kind of syntactic category, but that it *selects* only certain sorts of admissible complements. ‘Jon hopes Mary will get rich’ is perfectly okay, as is ‘Jon hopes for Mary to get rich’, and these allow us to say essentially the same thing. This shows that Broome is wrong about why ‘Jon ought Mary to get rich’ is ungrammatical. Like ‘Jon hopes Mary to get rich’, it is ungrammatical because of selectional properties of ‘ought’ and ‘hopes’, not because ‘ought’ is not a control verb.

I think this case is illustrative because Broome also believes that when I say that I don’t think the question of whether you stand in the OUGHT relation to the proposition that I exercise daily makes any sense, I am merely misled by the fact that the sentences we might use to state answers to that question are ungrammatical. Not only is ‘You ought me to exercise daily’ ungrammatical, so is Broome’s favored locution, ‘You ought that I exercise daily’. Broome (personal communication) has a hypothesis about why I find these claims incoherent—it is because they are ungrammatical.

That is an interesting hypothesis, but it overgeneralizes and yields false predictions. ‘Jon hopes Mary to get rich’ is ungrammatical, but it makes perfect sense. Its anomaly is syntactic, not semantic: I would know what anyone was saying who uttered it. It means that Jon hopes Mary will get rich. Similarly, ‘Jon ought that Jon gets rich’ is ungrammatical, but I don’t think it is incoherent either. It means, pretty obviously, that Jon



ought to get rich. What I don't understand is what it would be for it to be the case that Jon ought that Mary gets rich, unless it is supposed to be that Jon ought to make it the case that Mary gets rich, or to ensure that she gets rich—but those are both actions, things that Jon can do.

Now, in the face of this evidence, Broome (personal communication) still holds that I am being misled by grammar. He thinks that I am allowing that I understand an ungrammatical sentence only if I can paraphrase it with a grammatical sentence, and hence ruling out on grounds of grammar his view, which requires that there are some interesting claims that cannot be made with *any* grammatical sentence. I don't believe that I am making this error. I genuinely don't understand what it could possibly be for it to be the case that you ought that I exercise daily. I don't claim that this is an argument which will convince Broome, but I think that unless you have spent too much time poring over papers in deontic logic that treat 'ought' as taking propositions simply because it is formally tractable to model deontic logic on modal logic, you won't think that this is an intelligible question, either.<sup>32</sup>

Finally, I think that the fact that 'ought' belongs to a linguistic construction that can *so easily* take arbitrary subjects in its complement should draw our attention forcefully to the question of *why* there is *no* grammatical way of doing so in the case of 'ought'. *None* of the following sentences are grammatical with 'ought':

- 7f** Jon wants/\*hopes/\*ought Mary to get rich.
- 7g** Jon wants/hopes/\*ought for Mary to get rich.
- 7h** Jon \*wants/hopes/\*ought Mary will get rich.
- 7i** Jon \*wants/hopes/\*ought that Mary gets rich.

But each of these sentences is fine with either 'wants' or 'hopes', and **7g** is fine with both. Grant Broome that there are such interesting things to talk about. Once we observe that all it would take to be able to talk about them is to get around the kinds of minor differences that make the difference between which of **7f–7i** admit of 'wants' and which allow for 'hopes', it becomes extremely puzzling why we have never developed any of these means to be able to talk about them. I offer a simple hypothesis. We have never needed 'ought' to work in such constructions because there is no such interesting thing to talk about. As we saw in section 3.3,

32. Note that I am joined in this thought both by the father of deontic logic, G. H. von Wright (1951, 1981), and by one of its more prolific contributors, Hector Castañeda (see, for example, Castañeda 1981).

the same goes for a broader class of verbs like ‘neglects’, ‘forgets’, ‘proceeds’, and ‘aspires’.

### 3.5. The Basic Problem Again

The Basic Problem is not a syntactic problem. Some sentences, like ‘Jon hopes Mary to get rich’, are anomalous for purely syntactic reasons, making perfect semantic sense. The sentences required in order to relate you by the OUGHT relation to the proposition that I exercise daily are not like that. They are ungrammatical, true. But there also isn’t anything for them to be about. They are semantically anomalous. The problem that I have been claiming for the view that ‘ought’ relates agents to propositions is that it is too powerful. It makes sense of all of the right things, but then it also makes sense of things that don’t, in fact, make any sense. It is sensitive to *positive* evidence, making sense of all of the things that *should* make sense, but it ignores *negative* evidence, making predictions that some things *should* make sense that don’t.

Broome and Wedgwood are not put off by this consequence; rather, they have embraced it.<sup>33</sup> They both hold that we should be educated by this consequence of their theories—a discovery of a whole realm of interesting questions about whether you ought that I exercise daily, whether I ought that you brush your teeth, and so on. These are questions we could never have dreamed of asking before. My *reductio* is their interesting consequence, so I do not expect considerations like these to convince them. Nevertheless, this does not change the fact that, independently of theory, it is not what we would expect to find.

That is why I think that, in its relational sense, ‘ought’ relates agents not to propositions but to *actions*, in the very broad sense of things that agents can *do*. Actions, as I conceive of them, are not particular events but a kind of *property of agents*. On some, loose conceptions of properties, my exercising regularly may count as one of your properties. But being such that I exercise regularly is not something that you can *do*. So it is not an action. It is not the kind of thing that you can stand in the OUGHT relation to. Consequently, the view that ‘ought’ relates agents to actions correctly avoids overgenerating.<sup>34</sup>

33. Since I originally wrote this essay, Broome has moderated his stance on this point. See Broome n.d.

34. As Castañeda (1981) notes and Geach (1982) also suggests, somewhat less clearly, one payoff of this conclusion is that many versions of the Good Samaritan Paradox go away—even though I ought to help the injured person, and even though it

#### 4.1. Why It Matters

The questions I have been asking in this essay may seem very arcane and academic. After all, why should we care whether ‘ought’ relates agents to actions or to propositions? And how do all of these considerations about language really bear on philosophy or ethical theory? I think that any number of interesting questions in ethical theory can and do turn on the answer to this and similar questions. How can we get very far in trying to understand what we ought to do if we don’t even know whether the answer is supposed to consist in actions or propositions? And how can we evaluate proposals about the *semantics* of ‘ought’ sentences offered by expressivists and others as part of their metaethical views if we have no grasp of their structure?

To take a simple but important example, the end-relational semantics for ‘ought’ of Stephen Finlay (2009, 2010, forthcoming, n.d.) employs his assumptions about precisely these kinds of questions about ‘ought’ in order to tackle a wide range of the traditional questions of metaethical inquiry—and to give them *very elegant* answers if the underlying features of his view work out. If Finlay is right, then semantic analysis by itself can establish the truth of reductive normative realism in metaethics—quite a stunning and powerful conclusion. But Finlay’s view crucially relies on the premise that the raising ‘ought’ is the only ‘ought’.

In another important example, the expressivist semantics of Terry Horgan and Mark Timmons (2006) directly invokes the traditional view that ‘ought’ is a one-place predicate taking a propositional argument. According to Horgan and Timmons, while nonnormative sentences express what they call *is-beliefs*, ‘ought’ sentences express *ought-beliefs*—and *is-beliefs* and *ought-beliefs* are both psychological relations—albeit somewhat different ones—to propositions. It is essential to this view that *ought-belief* and *is-belief* are both relations to propositions, so it is essential to the view that ‘ought’ is essentially a context-independent propositional operator. In short, implications of the kinds of questions that I’ve been asking are fairly easy to come by. In the following sections, I outline just a few more examples of what I think are important potential implications of these questions.

---

is a consequence of my helping her that she is injured, it is not the case that I ought that she is injured—for that is not even the kind of thing that can be true. There are, however, other solutions to the Good Samaritan that are more general, so I don’t take this to be a chief virtue of the naive view.

## 4.2. Wide-Scoping

Broome is interested in whether ‘ought’ takes propositions partly because he is a prominent advocate of the ubiquitous *Wide-Scoping* program in the theory of rationality. In many different areas, it seems that someone’s situation can have an effect on what he or she ought to do or on what it is rational for him or her to do. For example, it seems that your ends, intentions, or desires, together with what is necessary for fulfilling them, can have an effect on what you ought to do. It seems that your intentions and your *beliefs* about how to fulfill them can have an effect on what you ought to do. It seems that your beliefs about what you *ought* to do can have an effect on what you ought to do, that your promises can have an effect on what you ought to do, and that your beliefs can have an effect on what you ought to believe. But in each of these domains, it seems too strong to suppose that it really *follows* from the fact that something is necessary for your ends (for example) that you ought to do it. On the contrary, putative counterexamples are common. You might have bad ends or irrational beliefs about what you ought to do.

Wide-Scopers propose to account for this by postulating ‘wide-scope’ *oughts* or requirements. The initial idea is that sentences like the following are scope ambiguous.

- 8a** If you will the end, you ought to take the means.
- 8b** If you believe that you ought to do it, then you ought to do it.
- 8c** If you believe that *p* and that if *p*, then *q*, then you ought to believe that *q*.

Since the ‘ought’ appears both in the consequent and in the sentence as a whole, it can be understood as taking scope either over the consequent or over the sentence as a whole. The former readings are the contentious *Consequent Scope* readings, which are subject to putative counterexamples; the latter are the *Wide Scope* readings, which are supposed to be not contentious. Wide-Scoping is commonly supposed to be “uncontroversial,” and rejecting it is supposed to be based on “confusion.”<sup>35</sup>

If **8a–c** and similar statements were uncontroversially true on some reading or other and were scope ambiguous, then the counter-

35. Compare prominent defenses of Wide-Scoping applied specifically to instrumental reason in Hill 1973 and Darwall 1983 and very broad advocacy of a Wide-Scoping approach in Gensler 1985 and Broome 1999. Also see the more restrained assessment of the consequences of wide-scope ‘oughts’ in Greenspan 1975 and the critical discussion in Schroeder 2004.

examples to their Consequent Scope readings would be evidence that Wide-Scoping is uncontroversial and of central importance in ethical theory. But if the deliberative ‘ought’ is a control verb, then there is no interesting question of whether it takes scope over the entire sentence or merely over its consequent.

It is possible, as I’ve noted elsewhere, to reconstruct Wide-Scope views as positive theories about various domains without the assumption that ‘ought’ takes propositions, so this doesn’t by any means show that Wide-Scope views are false. For example, the basic idea behind the Wide-Scope interpretation of **8a** can be expressed by saying that you ought to either will the means or not will the end. In the foregoing sentence, ‘either will the means or not will the end’ is a verb phrase, and it picks out a nontrivial property of agents. I think it describes something that you can *do*, in the very broad sense that I have in mind. So this evidence doesn’t show that Wide-Scope views are false. But it does undermine a major source of the idea that they should be *uncontroversial*, deriving from the idea that Wide-Scope readings are available as possible readings of conditional English sentences like **8a–c**, and hence from the idea that ‘ought’ takes propositions.<sup>36</sup>

#### 4.3. The Distinction between Agent-Relative and Agent-Neutral Reasons

It is plausible to suppose that whatever the deliberative ‘ought’ relates agents to is the same sort of thing as whatever *reasons* relate agents to. If ‘ought’ relates agents to actions, then reasons count in favor of actions. If ‘ought’ relates agents to propositions, then it is natural to think that, strictly speaking, reasons will count in favor of propositions. If this is so, the question of whether ‘ought’ takes propositions has important implications for some of our most treasured distinctions regarding reasons.<sup>37</sup>

The distinction between agent-relative and agent-neutral reasons is said by Thomas Hurka (2003, 628) to be “one of the greatest contributions of recent ethics.” So presumably this is supposed to be an important and uncontroversial distinction. But in fact, this distinction was successfully introduced to the literature only under the assumption that reasons ultimately count in favor of propositions. There is an

36. See Schroeder 2004 for further discussion.

37. See Schroeder 2007a and 2007b for further discussion of the issues raised in this section.

uncontroversial distinction in the neighborhood, but it is only with the aid of a controversial view such as the view that reasons relate agents to propositions that this uncontroversial distinction is taken to do the work that it is standardly assumed to do.

The uncontroversial distinction is a distinction between reasons that are reasons for any possible agent and reasons that are reasons for only some agents. This is where “free agent-variables” come in to the definition of “objective” and “subjective” reasons (the original words for ‘agent neutral’ and ‘agent relative’ in Nagel 1970). In order to get Nagel’s original formulation to work, we have to start, as he did, by *assuming* that all reasons are derivative from reasons to promote some state of affairs—so all *nonderivative* reasons really count in favor of a prospective state of affairs or proposition. The way that you check whether Ryan’s reason to promote some state of affairs  $p$  is agent-relative or agent-neutral is to look at the weakest relation,  $\lambda x, p(R)$ , which Ryan bears to  $p$ , such that necessarily, anyone who bears that relation to a state of affairs has a reason to promote that state of affairs. If ‘ $x$ ’ is free in ‘R’, then Ryan’s reason is agent relative; otherwise it is agent neutral. So given this definition, agent-neutral reasons are just reasons that are necessarily reasons for everyone if they are reasons for anyone at all—they are relations between agents and prospective states of affairs that everyone holds to that state of affairs if anyone does. Whereas by this definition, agent-relative reasons are just reasons that some agent can have, without other agents having the same reason. This is clearly an uncontroversial distinction.

But it is widely supposed that this uncontroversial distinction can help us to pick out the differences in the kinds of cases that can be allowed for by ordinary consequentialist views and those that can be allowed for only by views that look more like deontological views. But this turns on what sort of thing we think that reasons can stand in favor of. If reasons stand in favor of *propositions*, as Nagel assumed, then an agent-neutral reason for Franz not to murder will be a reason for every agent in favor of the proposition *that Franz doesn’t murder*. So on this view, if there are only agent-neutral reasons, and there is a reason for each agent not to murder, then given a choice between not murdering and murdering in order to prevent several other murders, there will be more reasons on the side of murdering than on the side of not murdering. And this is the ordinary consequentialist result about agent-centered constraints. So on the view that reasons stand in favor of *propositions*, the distinction picks out the relevant difference between consequentialist views and other moral theories.

But if reasons stand in favor of *actions* rather than propositions, and actions are, as I have been saying, simply a kind of nontrivial property of agents, we get no such result. On this view, if Franz's reason not to murder is an agent-neutral reason, then it is a reason for every agent in favor of the same thing as it is for Franz: *not murdering*. So on this view, if there are only agent-neutral reasons not to murder, nothing follows about the (im)possibility of agent-centered constraints. So if this view about reasons is right, then there is no uncontroversial distinction between agent-relative and agent-neutral reasons that has the implications that have been claimed for it and that have made it seem like "one of the greatest contributions of recent ethics."<sup>38</sup>

#### 4.4. The Viability of Deontology

Relatedly, one cause for suspicion about deontological moral theories has been that they seem to require agents to treat themselves specially—that they seem to require *different* things of each agent. Franz, according to this interpretation of deontological views, ought to avoid getting *Franz's* hands dirty, but Hans ought to do something else—to avoid getting *his* hands dirty. This is supposed to be a puzzling feature of deontological views—surely it would be simpler to think that moral obligations are the *same* for everyone rather than *different*—"neutral" rather than "relative."

The view that deontology requires different things of different agents is based on the view that what morality requires of agents is propositional in structure—that *Franz's hands not get dirty* may be required of Franz, but *it* is not required of Hans; what is required of Hans is that *Hans's hands not get dirty*. But on the view that what is required of agents is *actions*, this is not so. Actions, I have been saying, are a kind of property of agents, and so deontological views are naturally understood as holding that the *same* things are required of every agent—*not murdering*, *not stealing*, and the rest. So that is where the idea comes from that it is a peculiar feature of deontology that it requires different things of different agents, while consequentialism, sensibly, requires the same things of different agents. It is an artifact of a view about the kinds of thing that are the objects of

38. Convinced that there must be *some* good distinction between agent-relative and agent-neutral reasons that tracks issues related to the dispute between consequentialism and deontology, McNaughton and Rawling (1991) devised a quite different way of drawing the distinction. Their distinction is also based on controversial premises, however, albeit different ones; there is unfortunately no space to discuss the point here. See also Schroeder 2007a for further discussion.

*requirements*—the same family of views as the question of whether the agential ‘ought’ relates agents to actions or propositions.

#### 4.5. In Sum

These are just a few of the possible implications of the question of whether ‘ought’ takes propositions. In general, it is my view that we must get the answers to questions like this one right, or at least come to an adequate understanding of the costs of different answers, before we can hope to make progress on many other hard questions in ethical theory or even, as I argued in section 4.3, make the distinctions that we want to make in clear and uncontroversial ways. In this essay I have tried to assemble some of the relevant evidence in a way that allows us to make progress on the question. Though my arguments have at times turned on the presence or absence of relatively subtle readings of sentences, it’s my view that the evidence is strong that ‘ought’ has a deliberative, control sense as well as an evaluative, raising sense and the evidence is convincing that if this is right, then the deliberative sense ultimately relates agents to actions—understood as *properties of agents*—rather than to propositions. That interpretation is open to challenge; unlike some of my opponents, I don’t claim that those who disagree are merely subject to a linguistic confusion, and I haven’t by any means covered the full range of considerations either in linguistics or philosophy that bear on this question. But the important thing is that proponents of the propositional view must meet head-on the serious challenges for their view—particularly if they intend to marshal their view to advance some substantive normative or metaethical theory.

#### References

- Belnap, Nuel, Michael Perloff, and Ming Xu. 2001. *Facing the Future: Agents and Choices in Our Indeterminist World*. Oxford: Oxford University Press.
- Bhatt, Rajesh. 1998. “Obligation and Possession.” In *Proceedings of the UPenn/MIT Workshop on Argument Structure and Aspect*, ed. Heidi Hartley, 21–40. MITWPL.
- Broome, John. 1999. “Normative Requirements.” *Ratio* 12: 398–419.
- . n.d. “Reasoning.” Unpublished lectures given at Brown University and the University of Stockholm.
- Castañeda, Hector. 1981. “The Paradoxes of Deontic Logic: The Simplest Solution to All of Them in One Fell Swoop.” In *New Studies in Deontic Logic*, ed. Risto Hilpinen, 37–86. Dordrecht: Reidel.



- Chierchia, Gennaro. 1989. "Anaphora and Attitudes *De Se*." In *Language in Context*, ed. R. Bartsch, J. van Benthem, and P. van Emde Boas, 1–31. Dordrecht: Foris.
- Chisholm, Roderick. 1964. "The Ethics of Requirement." *American Philosophical Quarterly* 1: 147–53.
- Darwall, Stephen. 1983. *Impartial Reason*. Ithaca, NY: Cornell University Press.
- Finlay, Stephen. 2009. "Oughts and Ends." *Philosophical Studies* 143: 315–40.
- . 2010. "Normativity, Necessity, and Tense: A Recipe for Home-Baked Normativity." *Oxford Studies in Metaethics* 5: 57–85.
- . Forthcoming. "What *Ought* Probably Means, and Why You Can't Detach It." *Synthese*. Published online in *Springer Online First*, [www.springerlink.com/content/k43428281502n276/](http://www.springerlink.com/content/k43428281502n276/).
- . n.d. "A Confusion of Tongues." Unpublished book manuscript.
- Geach, P. T. 1982. "Whatever Happened to Deontic Logic?" *Philosophia* 11: 1–12.
- Gensler, Harry. 1985. "Ethical Consistency Principles." *Philosophical Quarterly* 35: 156–70.
- Greenspan, Patricia. 1975. "Conditional Oughts and Hypothetical Imperatives." *Journal of Philosophy* 72: 259–76.
- Grimshaw, Jane. 1990. *Argument Structure*. Cambridge, MA: MIT Press.
- Harman, Gilbert. 1973. Review of Roger Wertheimer, *The Significance of Sense: Meaning, Modality, and Morality* (Cornell University Press). *Philosophical Review* 82: 235–39.
- Hill, Thomas. 1973. "The Hypothetical Imperative." *Philosophical Review* 82: 429–50.
- Hilpinen, Risto, and Dagfinn Føllesdal. 1971. "Deontic Logic: An Introduction." In *Deontic Logic: Introductory and Systematic Readings*, ed. Risto Hilpinen, 1–35. Dordrecht: Reidel.
- Horgan, Terry, and Mark Timmons. 2006. "Cognitivist Expressivism." In *Metaethics after Moore*, ed. Terry Horgan and Mark Timmons, 255–98. Oxford: Oxford University Press.
- Horty, John F. 2001. *Agency and Deontic Logic*. Oxford: Oxford University Press.
- Horty, John F., and Nuel Belnap. 1996. "The Deliberative Stit: A Study of Action, Omission, Ability, and Obligation." *Journal of Philosophical Logic* 24: 583–644.
- Humberstone, Lloyd. 1971. "Two Sorts of 'Ought's'." *Analysis* 32: 8–11.
- Hurka, Thomas. 2003. "Moore in the Middle." *Ethics* 113: 599–628.
- Jackendoff, Ray. 1972. *Semantic Interpretation in Generative Grammar*. Cambridge, MA: MIT Press.
- Kratzer, Angelica. 1977. "What *Must* and *Can* Must and Can Mean." *Linguistics and Philosophy* 1: 337–55.
- . 1981. "The Notional Category of Modality." In *Words, Worlds, and Contexts*, ed. H. J. Eikmeyer and H. Rieser, 38–74. Berlin: Walter de Gruyter.
- Krogh, Christen, and Henning Herrestad. 1996. "Getting Personal: Some Notes on the Relationship between Personal and Impersonal Obligation." In

## Ought, Agents, and Actions

- Deontic Logic, Agency, and Normative Systems*, ed. Mark Brown and José Carmo, 134–53. New York: Springer-Verlag.
- Lewis, David. 1979. “De Dicto and De Se.” *Philosophical Review* 88: 513–43.
- MacFarlane, John, and Niko Kolodny. 2010. “Ifs and Oughts.” *Journal of Philosophy* 107: 115–43.
- McNamara, Paul. 2006. “Deontic Logic.” In *Stanford Encyclopedia of Philosophy (Spring 2006 Edition)*, ed. Edward N. Zalta, plato.stanford.edu/archives/spr2006/entries/logic-deontic.
- McNaughton, David, and Piers Rawling. 1991. “Agent-Relativity and the Doing-Happening Distinction.” *Philosophical Studies* 63: 167–85.
- Millsap, Ryan. n.d. “Practical Reasons and Means-End Transmission.” PhD diss. in progress, University of Maryland.
- Moore, G. E. 1922. *Philosophical Studies*. London: Kegan Paul.
- Nagel, Thomas. 1970. *The Possibility of Altruism*. Princeton: Princeton University Press.
- Nordlinger, Rachel, and Elizabeth Traugott. 1997. “Scope and the Development of Epistemic Modality.” *English Language and Linguistics* 2: 295–317.
- Radford, Andrew. 2004. *Minimalist Syntax: Exploring the Structure of English*. Cambridge: Cambridge University Press.
- Romero, Céline. 2005. “L’évolution syntaxique des verbes modaux dans l’histoire de l’anglais” (“The Syntactic Evolution of Modal Verbs in the History of English”). PhD diss., University of Paris 3—La Sorbonne Nouvelle, Paris. [www.celineromero.com/Evolution\\_syntaxique\\_modaux.pdf](http://www.celineromero.com/Evolution_syntaxique_modaux.pdf).
- Ross, Jake. 2010. “Personal and Impersonal Obligation.” *Journal of Philosophical Logic* 39: 307–23.
- Schroeder, Mark. 2004. “The Scope of Instrumental Reason.” *Philosophical Perspectives* 18: 337–64.
- . 2007a. “Teleology, Agent-Relative Value, and ‘Good’.” *Ethics* 117: 265–95.
- . 2007b. “Reasons and Agent-Neutrality.” *Philosophical Studies* 135: 279–306.
- Sidgwick, Henry. 1907. *The Methods of Ethics*, 7th ed. Indianapolis: Hackett.
- Thomson, Judith Jarvis. 2007. “Normativity.” *Oxford Studies in Metaethics* 2: 240–66.
- von Wright, G. H. 1951. “Deontic Logic.” *Mind* 60: 1–15.
- . 1981. “On the Logic of Norms and Actions.” In *New Studies in Deontic Logic*, ed. Risto Hilpinen, 3–36. Dordrecht: Reidel.
- Wedgwood, Ralph. 2006. “The Meaning of ‘Ought’.” *Oxford Studies in Metaethics* 1: 127–60.
- . 2007. *The Nature of Normativity*. Oxford: Oxford University Press.
- Williams, Bernard. 1981. “Ought and Obligation.” In *Moral Luck*, 114–23. Oxford: Oxford University Press.
- Williams, Edwin. 1994. *Thematic Structure in Syntax*. Cambridge, MA: MIT Press.