



Ought and agency

Daniel Skibra¹ 

Received: 13 January 2022 / Accepted: 18 July 2022
© The Author(s) 2022

Abstract

A thorny question surrounding the meaning of *ought* concerns a felt distinction between deontic uses of *ought* that seem to evaluate a state of affairs versus those that seem to describe a requirement or obligation to perform an action, as in (A) and (B), respectively. (A) There ought not be childhood death and disease. (B) You ought to keep that promise. Various accounts have been offered to explain the contrast between “agentive” and “non-agentive” *ought* sentences. One such account is the Agency-in-the-Prejacent theory (“AIP”), which traces the difference to a particular kind of ambiguity in the prejacent. This theory has been criticized as linguistically unviable. Indeed, I level a few novel complaints against AIP myself in the present paper. But AIP has a kernel of genuine insight which allows us to explain the contrast—that the distinction between agentive and non-agentive *ought* sentences owes in part to the way natural language encodes information about agency. I develop this idea into a novel account that, like AIP, traces the contrast to an ambiguity in the complement of the modal. However, according to the view I propose, the Coercion View, a linguistically-motivated coercion operation produces the necessary grammatical conditions for agentive *ought*, which in turn allow a kind of variadic function operator in the style of (in: Recanati, *Literal Meaning*. Cambridge University Press, 2004) to produce the semantic effect we see on display in agentive readings of *ought*. Having explained the mechanism by which we get this structure, I show that it corroborates some of the central intuitions underwriting agentive *ought*. I submit that the Coercion View offers an explanation of agentive *ought* to take at least as seriously as any of its competitors.

Keywords Ought · Deontic modals · Agency · Event semantics

✉ Daniel Skibra
danielskibra@gmail.com

¹ Department of Linguistics, University of Konstanz, Constance, Germany

1 Introduction

The word *ought* can be used to express a number of different senses (or “flavors,” as semanticists tend to say)—epistemic, buletic, teleological, deontic, and perhaps others besides. A thorny question surrounding the meaning of *ought* that has long attracted the attention of moral philosophers concerns a felt distinction between deontic uses of *ought*. On the one hand, there’s the *ought* that evaluates a state of affairs, and on the other, the *ought* that describes a requirement or obligation to perform an action.^{1,2} The names given to this phenomenon are diverse³ and philosophers sometimes train their attention on different features associated with the phenomenon. But the action-enjoining and state of affairs-evaluating distinction is pretty consistent across characterizations of the phenomenon, as the quotations below make clear. (The emphasis is in each mine.)

Harman (1973): “In one use *ought* represents a predicate of the **possible state of affairs...** In another use, *ought* represents a relation between **an agent and a possible course of action.**”

Geach (1982): “In the symbolism of von Wright’s original article ‘Deontic Logic’ the operators ‘O’ and ‘P’ for obligation and permission are attached, not to propositional letters, but to letters which stand in for general terms, and answer to **kinds of actions.** [...] [O]bligation essentially relates to an agent, it is somebody’s obligation; if instead we try to think of the ought-to-be-ness... of a situation involving the agent, then our thinking is going to be confused...”

Schroeder (2011): “[‘O]ught’ 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. [...] ‘ought’ also has an evaluative sense, on which it means, roughly, that were things ideal, some proposition would be the case.”

Examples tend to draw out the contrast vividly and evoke the felt distinction without much set-up. In line with the rough characterizations set out above, the sentences in (1) concern an agent’s performing of some action, and in (2) the evaluation of some state of affairs.⁴

¹ Such a distinction traces back at least to Sidgwick 1874, Prichard 1912, and is discussed in the early literature on deontic logic (e.g., in von Wright 1951) but for more recent discussion cf. Feldman 1986, Grice 2001, Horty 2001, Schroeder 2011, Broome 2013.

² For the view that the felt distinction is merely apparent, cf. Chisholm 1964, Williams 1981a, Ch 9.

³ Here’s a sample of the names given to a distinction in the conceptual neighborhood of the one drawn above: agentive vs. non-agentive (Chrisman, 2016), deliberative vs. evaluative (Williams, 1981a; Schroeder, 2011), relative vs. non-relative (Grice, 2001) ought-to-do vs ought-to-be (Feldman, 1986). This terminology doesn’t all track precisely the same distinction. Broome (2013)’s distinction between owned and un-owned *oughts* is also in the conceptual neighborhood, but for Broome, control over one’s action is not an essential ingredient of this interpretation of *ought*, so the link to agency is not as direct as in other philosophers’ conceptualization. So, there may very well be more than one distinction in this conceptual neighborhood. Cf. Humberstone 1991 for discussion of multiple ways for *ought* to be related to a subject or agent. And cf. Chrisman 2016, pp. 124–125 for a clear presentation of the distinction, in the context of Chrisman’s discussion of previous proposals to explain it.

⁴ The sentences are labeled according to the philosopher who supplied the example.

(1) Agentive:

- a. You ought to keep that promise. (Harman)
- b. John ought to beat up Tom. (Geach)
- c. Jay ought to give up smoking (Chrisman)
- d. Alison ought to get a sun hat. (Broome)

(2) Non-agentive:

- a. Tom ought to be beaten up by John. (Geach)
- b. Luckless Larry ought to win the lottery (Schroeder/ Chrisman)
- c. There ought not be childhood death and disease (Chrisman)
- d. Milton, you ought to be living at this hour. (Wedgwood/ Chrisman, channeling Wordsworth)
- e. The meeting ought to start at noon. (Schroeder)
- f. Alex ought to get a severe punishment (Broome)

Informants tend to recognize a difference in these examples, and philosophers easily generate them whether or not they think there is any deep cause of these interpretive differences.⁵

In the present paper, I will defend an underappreciated strategy for accounting for this contrast, showing that the strategy has some heretofore untapped resources for explaining the contrast, and arguing that these resources allow the strategy to evade the criticisms leveled against it in the literature. According to the view I'll be advocating, the contrast owes to a difference in the logical form between the (1)-sentences and the (2)-sentences (agentive and non-agentive *ought*-sentences, as I will call them), and that the difference is one traceable to the preajcent on which *ought* operates.⁶ I'll call the view I defend the Coercion View, because the difference in preajcent I'll be appealing to will hinge on the output of the semantic operation of *coercion*, which is a kind of repair-mechanism in semantic composition.⁷ If an operator's composition with an argument would result in meaning that is inconsistent, incoherent, or type-mismatched, coercion is the process whereby the meaning of the argument is reinterpreted (*coerced* into a new meaning) so that it can compose with the operator. It will be my contention that there is a particular kind of coercion that allows for the meaning we associate with the agentive

⁵ Williams (1981a), for example, denies that there is a distinction in these sentences, though he has no trouble producing examples showcasing the different interpretations. Though, as I understand it (from Broome 2012), Williams came to reject the conclusion drawn in this paper in a subsequent, but unpublished, lecture entitled "*Ought, must, and the needs of morality*".

⁶ In calling them "agentive", I purposely avoided the term "agential" in describing the *ought*-sentences I have in mind, even though this appears in the literature. (Cf. Finlay & Snedegar, 2014) This term is typically used to refer to the surface form of the (1)-sentences, differentiating sentences of the form 'S ought to ϕ ' where ϕ is some action, from sentences not of this form, like (2c). Though surface form will be important in what follows, what I want to denote are the deliberative interpretations of *ought*-sentences, agential or otherwise, so I'm employing the term 'agentive' for this. In doing so, I adopt the terminology used by Chrisman (2016) (eg., cf. p. 124), and intend to mean the same thing Schroeder (2011) means by 'deliberative *ought*'.

⁷ For discussion, cf. Moens and Steedman 1988, Jackendoff 1997, De Swart 1998, Pustejovsky 1998, Zucchi 1998, Koontz-Garboden 2007.

interpretation of *ought*. According to the Coercion View, the schema representing agentive *ought* is something like (3) (at a rather coarse level of granularity).

(3) OUGHT [$OP_{AGENT=x}(\phi)$]

$OP_{AGENT=x}$ is an operator which adds an agent argument to the clause, over and above whatever verbal arguments might occur in ϕ . Without further argument, this may look like a hopelessly *ad hoc* maneuver to explain agentive *ought*. But it's not an unprecedented proposal. The Coercion View posits a similar logical form to work in deontic logic and the logic of agency and proposes to model agentive *ought* “compositionally”, by means of stacking an obligation operator over an agency operator.⁸ In recent years, opponents of these views challenged that such a proposal was linguistically unmotivated. And if so, whatever kind of *ought* such a proposal models, it's not the agentive *ought* that we seem to use quite easily in natural language. The point of the present paper is argue that, not only is this decompositional approach not *ad hoc*, it is in fact well-motivated on linguistic grounds, integrates well with existing linguistic theories, and does a good job explaining agentive *ought*.

Here's how the proposal advanced by this paper fits in with the existing ones. Previous explanations of the agentive/ non-agentive contrast cluster around the following theoretical options. There are those philosophers who deny that there is a robust distinction and therefore don't feel the need to provide a substantial account of it.⁹ Among philosophers who think there is some robust difference between the agentive and non-agentive *ought*-sentences, some claim that this is due to a difference in the logical form underlying the sentences. Others attempt to maintain a uniform logical form for agentive and non-agentive *ought* sentences and appeal to some other mechanism to explain the contrast.¹⁰

The difference-in-logical-form approach itself comes in two variants, depending on whether it locates the difference in an ambiguity in the complement on which *ought* operates, or in the very meaning of *ought* itself. Consequently, I'll refer to these two explanatory strategies as the complement-ambiguity strategy and the *ought*-ambiguity strategy, respectively. The Coercion View is an instance of the complement-ambiguity strategy.

A prominent group of philosophers have embraced the *ought*-ambiguity strategy, including all of those quoted above. Geach (1982) proposes that agentive *ought* is a distinctive sense of *ought* which does not operate on propositions, but on actions.

⁸ Cf. Hilpinen 1973; 1974 for such an example. The *stit* logics of Belnap and colleagues are another. Cf., e.g. Belnap and Perloff 1988, Perloff 1991, Horty and Belnap 1995, Belnap et al. 2001

⁹ Williams was perhaps such a philosopher, though cf. fn 5 for evidence that he may have changed his position. Chisholm certainly was; cf. Chisholm 1964. Kratzer (1977, 1981, 1991, 2012) doesn't take a position on the debate, but Schroeder (2011) enlists the Kratzerian framework in this camp.

¹⁰ For example, Wedgwood (2006) posits a *sui generis* agent parameter in the index. Chrisman (2016) appeals to metasemantic considerations, drawing on Castaneda (1975)'s distinction between propositions and practitions. Finlay and Snedegar (2014) propose a contrastive account, where what makes an *ought* sentence agentive has to do with the set of alternatives an *ought* is evaluated against. These views maintain a uniform (non-ambiguous) conception of *ought*, and explain the agentivity of agentive *ought*-sentences without positing an ambiguity in logical form. (However, as a reviewer points out to me, Chrisman's view is actually consistent with a difference-in-logical-form approach at a high level of granularity, depending on how one spells out the difference between practitions and propositions, and how these map onto natural language clauses.)

Schroeder (2011) gives a more linguistically nuanced version of this strategy, where the distinction between agentive and non-agentive *ought* is said to correlate with some syntactic facts. Specifically, agentive *ought* is said to behave like a control verb, “controlling” the arguments of the subordinate verb, and non-agentive *ought* is said to behave like a raising verb, where the surface argument of *ought* is semantically only an argument of the embedded verb. Agentive *ought* therefore has two argument positions—one for an agent and one for an action (or action-*type* if you incline towards thinking of actions as dated particulars).¹¹

The most famous instance of the complement-ambiguity strategy is the Agency-in-the-Prejacent hypothesis (“AIP”). AIP’s explanation for the contrast gives a clear recipe for producing agentive *ought* sentences. However, the explanation given by AIP has come under fire in recent years, most prominently in Schroeder 2011, which points to some serious linguistic shortcomings of AIP’s explanation and recipe.¹²

Based on the characterization above, it should be clear that the Coercion View is similar to AIP in that both fall under the complement-ambiguity umbrella. Ultimately, the criticisms mooted against AIP’s explanation are decisive—agentive interpretations of *ought* sentences do not come about the way AIP proposes.¹³ But AIP’s particular “recipe” does not exhaust the complement-ambiguity strategy. Moreover, even the moribund version of AIP contains an insight worth retaining—that agentive *ought* has an important connection to the way that information about agency is encoded in natural language. I present the Coercion View as a superior version of the AIP; a version that not only evades the criticisms of the AIP, but is also better motivated by independent linguistic considerations.

There are a few other motivations for offering the Coercion View, aside from showing that there is a version of the complement-ambiguity strategy that has not been sufficiently appreciated in the literature. It’s this: the semantics for modals articulated by Kratzer has sometimes been taken to be committed to the view denying a robust difference between agentive and non-agentive *ought*.¹⁴ I’ve claimed that the Coercion View turns on some independently motivated ingredients in its explanation of the contrast. These motivations will partly come out of modal semantics, and so the Coercion View will show that a Kratzer-style semantics indeed supports the distinction, contra what is sometimes claimed in the literature. So, this paper will make a case that agentive *ought* is explainable in terms of its logical form. In arguing its case, it makes a primarily linguistic contribution to the debate. But this approach is necessary, since the literature on the topic has advanced linguistic reasons for eschewing the complement-ambiguity strategy, mainly because of the difficulties with AIP. I think a rehabilitation of the complement-ambiguity strategy is in order, and the Coercion View can provide it. This will have some logical and meta-ethical consequences, but exploring these consequences in detail will have to await future work. For now, I focus on showing the Coercion View provides a feasible account and is worth further exploration.

¹¹ For criticisms of Schroeder’s proposal that agentive *ought* has a control syntax, see e.g., Chrisman 2012, Lee 2021. In the semantics literature, cf. Portner 2009, pp. 187–188 for relevant discussion.

¹² Cf. also Chrisman 2016, pp. 115–117 for different criticism of AIP.

¹³ In fact, as will become clear, I think the problems with AIP run even deeper than has been appreciated. However, cf. Klimczyk 2017 for a defense of AIP from Schroeder’s criticisms.

¹⁴ Cf., e.g., Schroeder 2011, pp. 2–3, Chrisman 2016, Sect. 4.3.

The plan

In broad outline, the plan is to show how various grammatical resources conspire to produce a structure like (3) and to argue that this gives us a plausible account of agentive *ought*. Here's the plan in more detail: Section 2 describes some important features of agentive *ought*, which an account of agentive *ought* should be in a position to explain. Section 3 describes AIP and its shortcomings. Though it levels some novel criticisms against AIP, it aims to point out the kernel of insight the Coercion View picks up from AIP. Section 4 lays out the Coercion View. This section will be concerned to show that the schema laid out in (3) can be derived from independently motivated mechanisms that have little to do with agentive *ought*. Since it turns out that 'OP_{AGENT=x}' is actually the result of multiple operations acting in concert, this section explains the "ingredients" of the view in detail. Section 5 answers some questions about the Coercion View, showing that the details yield predictions that accord with natural interpretations of agentive *ought* sentences. Section 6 shows how the Coercion View explains the features of agentive *ought* described in Sect. 2, and how it evades the criticisms of the other complement-ambiguity strategy, AIP, laid out in Sect. 3. Section 7 concludes. Although I'll be postponing discussion of the meta-ethical consequences of the Coercion View for another occasion, I will still provide a few words for why meta-ethicists should care about the Coercion View, about *ought*'s being unambiguous, and about the significance of a viable complement ambiguity strategy.

2 Hallmarks of agentive *ought* sentences

Aside from the intuitive distinction appealed to in the contrast between (1) and (2), why think there is any kind of special difference between agentive and non-agentive *oughts*? This section will provide some hallmarks of agentive *ought*. Some of these hallmarks have been offered as further evidence of the distinction.

2.1 Agency sensitivity

Philosophers have pointed out that there's an intuitive difference between S's performing an action and S merely being a participant in an event where an action-like behavior was performed. Chrisman (2016) provides a clear discussion of this point. Consider the sentence in (4) below. A merely "action-like" behavior is performed if the event is one we'd unreflectively think of as a kind of action (like kissing), provided one's performance isn't a genuine exercise of one's agency.

(4) Tom ought to kiss Bill.

As pointed out in Chrisman 2016, p. 115, there's an intuitive difference between (i) someone bringing it about that Tom kisses Bill by spiking his drink so that it foreseeably leads to his kissing Bill in his drunken haze, and (ii) Tom's just performing the action of kissing Bill "directly and actively". The difference can be further illustrated by considering the role of deviant causal chains. We might say that Tom's proper exercise of his agency involves forming an intention and having this intention be the cause of his

kissing Bill. Suppose that Tom resolves to kiss Bill, and that his having settled on this intention makes him so nervous that he stumbles forward. Rather than leaning in and deliberately planting his lips on Bill's, he lands mouth-first onto Bill. In this deviant causal chain, Tom winds up obtaining the intended result—his kissing Tom—but this outcome comes about through some “deviant” means rather than the paradigmatic way an agent's intention leads to action. An account of agentive *ought* needs to be agency sensitive in that it could in principle distinguish the (i) cases from the (ii) cases.¹⁵ In other words, it should distinguish between when the *ought* sentence requires genuine action on the part of the subject and when an outcome of merely action-like behavior (because deviantly caused, for example) is sufficient.

2.2 Event role asymmetry

Examples like the following have been given by Harman (1973) and Geach (1982) (with Geach attributing the example to Anselm). Suppose Tom and Bill are together attending a dance of some sort. Tom has been neglecting Bill all night, having failed to dance with him when he promised to do so. He really owes Bill an apology, and should finally dance with him while he still has the opportunity. Bill, by contrast, has suffered this indignation quietly—he has done nothing wrong. Now consider (5a) and (5b).

- (5) a. Tom ought to dance with Bill.
b. Bill ought to dance with Tom.

Under the circumstances just described, it seems to many people that (5a) and (5b) are subtly different. (5a) is true, but (5b) doesn't seem to be. Tom owes Bill a dance; Bill doesn't owe Tom anything (except perhaps an earful). To say he's got anything like an obligation to dance with Tom under the circumstances seems perverse. By contrast, Tom owes it to Bill to dance with him.

The *dance-with* relation is symmetrical, and the state of affairs where Tom dances with Bill is co-extensive with the state of affairs where Bill dances with Tom. So if *ought* just evaluates states of affairs, the thought goes, there's no reason to treat (5a) and (5b) as any different. But if *ought* can single out the agent of the described action as the one to whom the obligation adheres, we get a purchase on why we're inclined to think of (5a) and (5b) differently.

Similar considerations apply to passive transformations of sentences. As with the symmetrical relations above, the state of affairs where Tom kisses Bill is co-extensive with the state of affairs where Bill is kissed by Tom. Yet, as above, it seems to many people that (4) and its passive transformation—*Bill ought to be kissed by Tom*—are subtly different. Specifically, people report the intuition that (4) can be true, without

¹⁵ As a reviewer observes, it may be that the meaning of the verb *kiss* requires a connection between intention and success that would render *Tom kissed Bill* false in such deviant causal chain circumstances. But there are other verbs where the connection isn't so tight. Consider the verbs *shoot* or *break* (on its transitive reading). These verbs both have readings where the deviant causal chain circumstances would still allow the sentences *Tom broke the window* or *Tom shot the target* to be true. When these sentences are embedded under *ought*, they'd provide even more stark examples where our account of agentive *ought* needs to be agency sensitive. Cf. Pietroski 1998, Sect. 5, for relevant action-theoretic discussion, though the discussion doesn't touch on considerations concerning *ought*.

its passive transformation being true. Again, for this to be the case, it seems that in such an interpretation, (4) is not just evaluating the state of affairs, but singling out the agent of the described action.

I've called this "event role asymmetry" because it suggests that the relevant interpretation of *ought* can single out a participant in the described event as playing a distinctive role. An account of agentive *ought* should be able to explain this role asymmetry, even in states of affairs where the relation described by the verb is symmetrical.

2.3 Future orientation

Now to state an observation that might be obvious: agentive *ought* sentences typically, even paradigmatically, contain verbs of action. This observation has some less obvious consequences. Action verbs are categorized as *eventive* verbs. These are verbs that describe a kind of event taking place, as opposed to a state obtaining. Insofar as they contain action verbs, agentive *ought*-sentences have eventive complements. And modal sentences with eventive complements have a distinctive temporal profile—they are all *future oriented*. This terminology comes from Condoravdi (2002), who distinguishes between modals' temporal *orientation* and temporal *perspective*. In terms of the deontic modals that are the focus of this paper, the temporal perspective is the time relative to which the obligation-relevant circumstances obtain. The temporal orientation has to do with the temporal location of potentially true-making events described by the prejacent. Compare (6a), with the eventive *go to his office*, with (6b), with the stative *be in his office*, below.

- (6) a. Jay ought to go to his office.
b. Jay ought to be in his office.

The temporal perspective of both of these is present, since the circumstances that call for Jay to *go to his office* or *be in his office* are the ones obtaining presently. But they differ in their temporal orientation. (6a) is future oriented, because the time at which Jay's going to the office is evaluated extends into the future. An intuitive way of thinking of this is that his going to the office at some unspecified point in the future would count as complying with the conditions set by the prejacent. (6b), on the other hand, is present oriented, since his being at the office or not is evaluated presently. (His showing up at his office half an hour from now would not meet the condition described by the prejacent.)¹⁶

The examples of agentive *ought* sentences provided in (1) are all future oriented. And this is no surprise, because they all contain eventive verbs. The non-agentive examples in (2) are more varied. Some are future oriented, some are present oriented. So, non-agentive *ought* sentences can be either future or present-oriented. There's an important linguistic generalization lurking here, which is that agentive *ought* sentences are a subset of the future oriented *ought* sentences. Insofar as this generalization claims that *all* agentive *ought* sentences are *future oriented*, it might give you pause. What about *Jay ought to have gone to his office*? In fact, if this is agentive, it is no

¹⁶ Since (6b) contains the stative *be in his office*, present orientation is the default interpretation, but it may have a future orientation—as in *Jay ought to be in his office tomorrow afternoon*. For (6a), the future orientation is obligatory; there is no present oriented interpretation.

counterexample to the generalization. The distinction between temporal perspective and orientation shows us why. It is the temporal perspective of the modal that is shifted into the past, due to the Perfect.¹⁷ But the orientation is still future. So, at some point in the past, Jay's circumstances (then) were such that his (future) going to the office (relative to those circumstances) was called for.

I mentioned *paradigmatic* agentive *ought* sentences, which contain verbs of action. What about putatively agentive *ought*-sentences with stative prejacent, like *Ghengis ought to live a life of peace* or *Martha ought to be happy*? Well, modals with stative complements are typically present oriented, but have a permissible future oriented interpretation.¹⁸ My contention is that when such sentences have an agentive interpretation, they are *ipso facto* future oriented. I won't belabor this particular point, but here's a quick argument in favor of it. Agentive interpretations of *ought* concern putative actions that an agent can take. Suppose *Martha ought to be happy* were present oriented. Then, if she were not presently happy, nothing she could do would allow her to comply with this advice. The only way for her to meet the condition is for her to change. Then she must bring it about that she is happy, and given that she is presently *not* happy, it can only be made true in the future, if at all. So, in order for this to have an agentive interpretation, it must be future oriented. Though this is just a sketch of an argument, it comports with some well known behavior of stative predicates—they admit of a change-of-state interpretation in certain contexts (an *inchoative* interpretation). And when these are in the scope of a modal, they result in a future orientation.

There's one more case worth addressing. Can you have agentive readings of *ought* sentences with progressive prejacent, like (7)?

(7) Right now, Dylan ought to be doing his homework—not tomorrow, not in five minutes, right now!

The progressive (*be doing...*) is thought to denote the ongoing state of an event unfolding or taking place, and *now* is traditionally thought of as an indexical denoting the utterance time.¹⁹ So if (7) can have an agentive interpretation and yet be felicitous with the temporal adverbial *now*, it would make problems for the generalization offered here. While there is a prominent non-agentive interpretation of (7), I don't want to deny that it has an agentive interpretation as well. But I do think the agentive interpretation is still future-oriented. What allows me to say this is that, although *now* is usually thought to have pick out the utterance time, this a bit of an oversimplification, since there are attested discourse-bound uses where it functions to denote the consequent state of a prominent event.²⁰ Where the prominent event is the utterance event, the resulting interpretation would be genuinely future-oriented and comport with the earlier remarks about inchoative interpretations of stative predicates. (While *do one's homework* is eventive, *be doing one's homework* is a derived state.)²¹

¹⁷ Cf., e.g., Condoravdi 2002, pp. 75–77 for an analysis. It's the *have* + past participle that is the surface form of the Perfect.

¹⁸ Cf. the point in fn. 16.

¹⁹ Cf. e.g., Parsons 1990, Ch. 9 for relevant remarks on the progressive, and Kaplan 1977 for this traditional interpretation of *now*.

²⁰ Cf., e.g., Stojnić and Altshuler 2021.

²¹ According to my remarks here, what the agentive interpretation of (7) says is that the consequent state of the utterance time is one where an event of Dylan's homework-doing is ongoing. This is future-oriented,

In sum, the generalization that agentive *ought* sentences are future oriented is sound. To be clear, the key observation of this section isn't that agentive *ought* sentences are future oriented because they all contain action verbs. That wouldn't be true. It's that they usually contain verbs of action, and these are future oriented. When an *ought* sentence contains a stative verb in its complement, the agentive interpretation of the sentence is one which has a future orientation. In light of this descriptive generalization, an account of agentive *ought* should be able to explain the relation between agentivity and future orientation.

3 Agency-in-the-prejacent theory

In this section, we'll see how AIP works and why it fails. In addition to summarizing some of the extant criticisms of AIP, I will add some of my own. I do this because I think that AIP contains a good insight—that the agentivity of the targeted reading of the *ought* sentences is a reflex of the way language encodes information about agents. It just exploits this insight in the wrong way. So my motivation in this section is to make explicit exactly how the proposal fails to provide a satisfactory explanation of the distinction in order to save the legitimate insight.

AIP appeals to the linguistic proposal that sentences have a way of encoding information about the agent of an action (and about participants of events more generally) via the lexical semantics of the verbs contained in the sentence. The explanation for agentive *ought* given by AIP draws on the theory of thematic relations and their corresponding theta roles. It attempts to capture the notion that some *ought*-sentences of the form *S ought to ϕ* variably have agentive readings by appealing to an ambiguity in the prejacent, *S ϕ s*. The sense of agency owes to the fact that the sentence itself contains information about the agency of the participants in the event described by the Verb Phrase.

In verbal semantics, verbs stand in *thematic relations* to their arguments—the verb's arguments encode information pertaining to the object's role in the event or the action described by the verb. This has a syntactic reflex as well, since a given verb has a specified number of arguments which encode information about thematic relations (“theta-roles”). A verb like *kiss* has two argument positions—one expressing the agent relation and one expressing the theme relation, where the former is the agent performing the action described by the verb and the latter is the object being acted upon in the event described by the verb. Of course, there are other thematic relations one could express linguistically as well by means of optional *adjuncts*.²² The *arguments* of a verb are of a distinguished sort. A sentence missing an argument is typically ungrammatical because the verb “assigns” a certain number of theta-roles which need

Footnote 21 continued

and conforms to the generalization. That there is such an interpretation with *now* is bolstered by the fact that *now* is no less felicitous with an eventive complement which is uncontroversially future oriented: *Dylan ought to do his homework now*. But these are subtle issues, made more complicated by the progressive, and deserve some more attention. I thank a reviewer for raising this interesting question.

²² You could, for example, provide the *location* of the kissing event. Also, *instruments* are a thematic relation most often expressed by means of adjuncts. You could indeed say *Tom kissed Bill with his lips*, where *his lips* would express the instrument thematic relation.

to filled.²³ For a verb like *kiss*, the theme argument is obligatory—*Tom kissed* is ungrammatical.^{24, 25}

Let's reconsider an example discussed earlier, repeated here as (8), to check the kind of predictions made by AIP.

(8) Tom ought to kiss Bill.

According to AIP, (8)'s agentive interpretation arises because, first of all, *ought* acts as a sentential modal operator along these lines: 「*OUGHT* [Tom kiss Bill]」. Secondly, the clause embedded under *OUGHT* has an argument that is theta-marked as *AGENT* by the verb *kiss*. To make this explicit, (8) is then analyzed as (9).

(9) Ought [Tom_{AGENT} kiss Bill_{THEME}.]

On this view, there's nothing in the logical syntax of *ought* responsible for the agentive interpretation. Rather, the agentive interpretation owes to semantic information contained in the prejacent itself. Since the prejacent in (9) is one which contains an *AGENT* argument, the *ought*-sentence of which it is part will be agentive. It licenses a mapping to a logical schema which makes the agency explicit. Since AIP was proposed by proponents of *stit* logics, the natural language sentence is then mapped to a deontic modal operator and *stit* operator. But what concerns us is the motivation for this mapping given by AIP; we needn't concern ourselves with the details of *stit* logics.

AIP does a good job dealing with some of the explananda discussed earlier. For example, it builds a robust sense of event role asymmetry right into the explanation of the phenomenon. A sentence will at most have one agent argument, even when the verb describes a symmetrical relation.²⁶ So, *dance* or *kiss* will have an argument specified (or theta-marked) as the *AGENT*, and another theta-marked with some other theta-role. It's more difficult to say how AIP handles agency sensitivity. According to AIP, (8) will be interpreted as agentive in virtue of *Tom* occupying the *AGENT* theta-role. But will the truth-conditions for (8) be satisfied if Tom's kissing Bill comes about through some deviant causal chain? A proponent of AIP will have to deny that they are if AIP is to distinguish between Tom's genuine agency and his mere participation in action-like behavior, though there's no evidence that the theory of theta roles and thematic relations is sensitive to this.

Probing a bit further, we find that AIP's recipe for detecting agentive *ought* massively overgenerates. Schroeder (2011) presses these objections most forcefully.

²³ This is the significance of subscripting an argument of the verb with, say, *AGENT*. Such subscripting indicates which argument is assigned the agent theta-role by the verb.

²⁴ I say "typically" because there are interesting cases here. *AGENT* theta arguments are obligatory; *Kissed Bill* is ungrammatical. But passive transformations will often allow agent arguments to go unexpressed, as in *Bill was kissed*, though it has been argued by some syntacticians that the agent argument is in fact still present, even if only implicitly. Cf. Roeper 1987 or the discussion in Glanzberg 2009.

²⁵ *Theme* is sometimes used interchangeably with *patient*, though some theorists differentiate the two on the basis of whether or not the thing acted upon changes its state as a result of the event it undergoes. If we opt for this finer-grained distinction of *theme* versus *patient*, in the sentence *Tom cut the cake*, *the cake* would be the patient, and not the theme, because it's changed its state as a result of the cutting event.

²⁶ NB: Of course, the *AGENT* argument could be plural, and denote multiple agents, but the point still stands.

Consider that passive transformations of *ought* sentences tend not to have agentive interpretations. So, *Bill ought to be kissed by Tom* will typically have a non-agentive interpretation. But the sentence has an agent argument all the same, so AIP predicts that it would be agentive. Schroeder argues that the problem runs even deeper—AIP both systematically over- and undergenerates. He gives the following example to illustrate.

- (10) a. Luckless Larry ought to win the lottery.
 b. It ought to be the case that Luckless Larry wins the lottery.

According to Schroeder, the natural reading of (10a) is non-agentive—it was one of the paradigm examples of a non-agentive *ought*-sentence in (2). Yet, as Schroeder (2011) claims, there is a remote, but still accessible reading where (10a) is indeed agentive—think of a context where Larry is able to fix the lottery and ensure that he wins.²⁷ A verb like *win* does not have an AGENT argument, but only, say, an EXPERIENCER argument. (EXPERIENCERS are typically thought of as the theta-role appropriate for subjects of eventualities requiring sentience of some sort, but which don't involve agency, like psychological state predicates.) AIP would then undergenerate because if there is such an agentive reading of (10a), however remote, the AIP doesn't predict it.

AIP has a bit of latitude here. It could explain the availability of the agentive reading of (10a) by appealing to an *ambiguity* in the prejacent, where the prejacent itself has a remote agentive reading which then makes the agentive reading of the *ought*-sentence available. On this hypothetical agentive reading, the verb *win* takes an argument which is theta-marked as AGENT, whereas on the non-agentive reading, it takes an argument not theta-marked as AGENT but EXPERIENCER instead. But Schroeder thinks this putative fix runs us back into the overgeneration problem. Since (10a) and (10b) both presumably have the same prejacent, namely that *Luckless Larry wins the lottery*, (10b) should have an agentive reading as well, but Schroeder denies that such a reading exists. If not, this is troublesome for AIP because there is no apparent difference in the prejacent between (10a) and (10b)—why should there be an ambiguity in the embedded clause in (10a), but not in (10b)?

There is some dispute in the literature about whether Schroeder is right that (10b) can never have an agentive interpretation. For example, Chrisman (2016) and Bronfman and Dowell (2018) deny that (10b) is incapable of receiving an agentive reading, and I'm inclined to agree with them.²⁸ If (10b) can have an agentive interpretation in some contexts, then it would seem that Schroeder can't use this example to underwrite this particular overgeneration claim. Likewise, the proponent of AIP can appeal to the same kind of ambiguity in the prejacent to say why *Bill ought to be kissed by Tom* can lack an agentive reading.

The bigger problem for AIP has less to do with its ability to track available agentive readings than with the mechanism it posits to do so. The kind of ambiguity AIP needs to appeal to amounts to a verbal polysemy claim. When the *ought*-sentence has an

²⁷ Actually, Schroeder supposes that on the remote agentive reading, (10a) is just false. But Chrisman (2016) discusses conditions, like the ones just suggested, under which it might be true.

²⁸ The interpretive issues here are rather delicate, and I discuss them in Sect. 6, so I will defer further discussion of this question until then.

agentive reading, the embedded verb selects for an AGENT argument and when it does not have an agentive reading, it does not. So the answer buys modal uniformity at the cost of committing to verbal polysemy. This may seem reasonable – verbal polysemy isn't particularly uncommon. The problem is that *this is not how verbal polysemy works*. AIP's plausibility hinged on its appeal to a theory whereby verbs encoded information about agency via their relevant theta arguments. But it is not part of that theory of verbal argument structure that verbs can be ambiguous with respect to whether or not their external argument is an agent or something else.

Why not? Two reasons strike me as decisive on this point. First, foreshadowing next section's discussion, agent arguments are *external arguments*. So *Luckless Larry* would be an external argument for *win* whether it has the natural EXPERIENCER interpretation or the exceptional AGENT interpretation AIP proposes. External arguments are distinguished from *internal* arguments in part because they tend *not* to influence a shift in meaning of the verb.²⁹ AIP's polysemy claim amounts to proposing that verbal meaning can be affected by the external argument, but the linguistic evidence points in the opposite direction.

Secondly, the kind of ambiguity for *win* that AIP posits only seems to exist in the scope of the appropriate modal. Let's say that *win** is the verb just like *win* but that differs in its thematic grid. It has an AGENT argument instead of an EXPERIENCER argument. What AIP needs to say is that *win** is the word that occurs in the agentive version of (10a). This means that AIP effectively posits a distinct meaning expressed by a new lexical item *win** that we simply don't detect in unembedded positions, but only under *ought*. This is not a good sign for AIP's proposal. If *win** only occurs in very specific embedding conditions, the hypothesis to beat would be that the phenomenon we are trying to explain owes to the embedding environment, not to lexical polysemy, which we should expect to observe unembedded.³⁰ So, the polysemy claim AIP would posit to deal with these challenges is a nonstarter, from a linguistic standpoint. But without the polysemy claim, AIP is powerless to deal with Schroeder's over- and undergeneration objections. AIP would then be unable to explain the presence of an agentive interpretation for (10a), just as it would be unable to explain why *Bill ought to be kissed by Tom* lacks one. Nonetheless, there is a kernel of insight in AIP I think should be retained. AIP attempts to marry the agentivity of *ought* to a linguistic theory about how agency is encoded in language. This seems like a promising and well-motivated strategy. However, in explaining agentive *ought* by appealing to an implausible kind of verbal polysemy, AIP runs afoul of some of the data that underwrites this theory.

²⁹ For example, *win the race*, *win at life*, and *win his heart* are plausibly thought to involve different but related senses of *win*. It is precisely the internal arguments (*the race*, *at life*, and *his heart* respectively) that trigger the different senses at issue here. The external argument tends *not* to affect the meaning of the verb in this way. Cf. Bresnan 1982, Grimshaw 1990 and especially Marantz 1984 for this observation. I will expand on this point in the next section.

³⁰ Schroeder (2011, p. 13) makes a similar objection when he points out that we don't find this kind of ambiguity under other raising verbs like *to seem*. My objection here is related, but slightly different. My contention is that, if the ambiguity were really due to two *verbs* with different meanings (even if the two are related), then we would expect to find *unembedded* instances of sentences with the *win** meaning, but we don't. Thanks to a reviewer for suggesting I highlight the difference.

4 The Coercion View

The Coercion View is similar to AIP in that it locates the difference between agentive and non-agentive *ought* in an ambiguity in the complement of the modal. But the mechanism which yields the agentive reading is quite different from AIP. Let me start with an informal first pass at the Coercion View, by considering what Recanati (2004, pp. 107–109) calls variadic functions. These are functions from relations to relations which increase (or decrease) the adicity of the relation. So, adding a predicate modifier to a predicate expressing an n -ary relation results in a predicate expressing an $n + 1$ -ary (or $n - 1$ -ary) relation. Recanati calls upon variadic functions to explain a certain kind of phenomenon, where the truth conditional content of a sentence includes components that aren't overtly specified.³¹ For example, consider sentence (11).

(11) It's raining.

In (11), no location is specified, yet hearers typically understand the sentence as concerning rain *at a specific location*—the location of the utterer, say. To get the sentence to reflect the truth conditions which include the location, the variadic function operator $OP_{location}$ will map the RAINING relation to the RAINING-AT relation by increasing the adicity of the RAINING relation by 1, thereby providing an argument position for a location (where l is the location of the raining event).

(12) $OP_{location}(\text{Raining}) = \text{Raining-at}(l)$,

The Coercion View makes a similar claim about the source of agentive interpretations of *ought*. According to the Coercion View, the difference between agentive and non-agentive *ought* sentences is that former have a kind of operator which allows for an additional argument position to be added to the clause which is related to the modal in a distinguished way. Consider (13).

(13) OUGHT [$OP_{AGENT=x}(\phi)$]

Notice that in (13), OUGHT itself is not ambiguous. OP_{AGENT} is a variadic function which takes the 0-place relation ϕ and turns it into a 1-place relation that can take an agent argument. Because the source of the agentive interpretation lies with a kind of ambiguity in the modal's complement, the proposal is a form of the complement-ambiguity strategy. And if we consider the additional operator OP_{AGENT} part of the prejacent, we have a variation on the AIP idea. But it is different from the erstwhile AIP explanation in that the source of agentivity is the additional operator, *not* the thematic grid of the verb in ϕ , which remains unchanged.

There's also an important difference between the kind of operator I'm positing here, and Recanati's variadic functions. For Recanati, they are thought to rely on a process of pragmatic enrichment, whereby the linguistic expressions associated with our utterances are pragmatically enriched with supplemental material before they are truth-conditionally evaluated.³² Proponents of pragmatic enrichment usually appeal to various sorts of evidence for the existence of these kinds of pragmatic processes.

³¹ Readers familiar with the literature about the semantics-pragmatics interface will recognize this as part of the debate surrounding unarticulated constituents. (Cf. Perry 1986, Recanati 2002.)

³² Cf. e.g., Hall 2008, Recanati 2010 on pragmatic enrichment.

My proposal doesn't appeal to pragmatic enrichment, but to *coercion*. I will argue that OP_{AGENT} is actually *grammatically* motivated by several factors relevant to the interpretation of the sentence wherein it occurs. The proposal therefore has more in common with the kind of contextualism proposed by Stanley (2000), Stanley and Szabo (2000) than with the enrichment accounts of Recanati. To properly explain the Coercion View, then, it will be necessary to explain the nature and source of this coercion, which I will do in the following two subsections. Importantly, what I'm schematizing here as an operator, OP_{AGENT} , is actually the joint operation of two different components. The first has to do with how agency is encoded in language, as promised in the last section. The second has to do with the interpretation of modals. I explain these two components in turn.

4.1 Encoding agency in language

The points to emphasize from this section are as follows. First, agent arguments are "external arguments", which compose with the Verb Phrase via a special composition operation called Event Identification. Second, the kind of external argument that can so-compose is constrained by properties of the event described by the verb and its internal arguments. Explaining what this means will require some unpacking.

External arguments are those which are the "farthest" from the verb. This is a sketchy notion which is cashed out differently in different theories, but the common thread amongst them is that they are the last argument positions of the verb to be saturated, which is why those others are more "internal".³³ It is tempting to say that external arguments are the *subjects* of sentences. This is on the right track, but it puts the cart before the horse, syntactically. Sentences in active voice are the unmarked case, and unmarked sentences have the external object (if the verb has one) as the sentential subject.

There's quite a bit of syntactic evidence for the claim that external arguments are "farther" from the verb, and this evidence usually figures prominently in theories about the behavior of active, passive, and middle voice. But there's also semantic evidence that differentiates internal from external arguments that's particularly salient for us, and this is the observation that the sense of the verb depends on the verb's internal and not external arguments. The data this generalization is based on is quite striking. There are many cases, as shown by the following examples, due to Kratzer (1996) (and drawing on an argument by Marantz (1984)), where an internal argument triggers a particular interpretation of the verb.³⁴

- (14) a. throw a baseball
 b. throw support behind a candidate
 c. throw a boxing match
 d. throw a party
 e. throw a fit
- (15) a. kill a cockroach

³³ Cf. Williams 1981b, Chomsky 1993 for discussion.

³⁴ Cf. also Bresnan 1982, Grimshaw 1990 for this observation, though they ultimately explain it in a different way than Marantz and Kratzer.

- b. kill a conversation
- c. kill an evening watching TV
- d. kill a bottle (i.e., empty it)
- e. kill an audience (i.e., wow them)

The thing to notice here is that the meaning of the verbs—*throw* and *kill*, above—varies depending on the nature of the internal arguments. The fact that *throw* in *throw a baseball* has the meaning it does has to do with the fact that the event described by *throw a baseball*—one where the baseball’s location is displaced by means of a launching movement—is the kind of thing that can be done with baseballs, whereas the kind of meaning it has in *throw a party* (organizing and hosting an event) has to do with the fact that such a throwing-event is the kind of thing that can be done with parties. There is an asymmetry here, because you won’t get the *throw a party*-type meaning from the choice of an external argument.³⁵

On this basis on these and similar kinds of data, Kratzer proposes to “sever” the external argument from the verb. In the simplest sense, severing just means that the external argument is no longer considered part of the lexical meaning of the verb. To illustrate part of what’s at stake here, we can compare Kratzer’s proposal to two other prominent accounts of verbal semantics, using Parsons (1990)’s stock example, *Brutus stabbed Caesar*. Note that the schematizations below don’t represent the sentence as a whole, but just what each proposal takes to be the semantic contribution of the verb *stab* to be.

- (16) a. Brutus stabbed Caesar.
 b. Davidson: $\lambda y.\lambda x.\lambda e[\text{STABBING}(e,x,y)]$
 c. Neo-Davidsonians: $\lambda y.\lambda x.\lambda e[\text{STABBING}(e) \ \& \ \text{AGENT}(e,x) \ \& \ \text{THEME}(e,y)]$
 d. Kratzer: $\lambda y.\lambda e[\text{STABBING}(e,y)]$

In (16b), based on Davidson (1967)’s proposal for action sentences, the arity of the predicate STABBING contains argument “slots” for an event variable, as well as the verb’s (internal and external) arguments.³⁶ In (16c), the Neo-Davidsonian approach associated with Higginbotham 1985; Parsons 1990, STABBING has only a slot for an event argument, and the internal and external arguments are conjoined via predicates denoting their thematic role. Kratzer’s proposal is given in (16d), where STABBING has slots only for an event argument and the internal argument(s).

Were this *just* a proposal about verbal meaning, it might not be too important for our consideration of agentive *ought*. But this is also a proposal about how external

³⁵ Let me be clear about this point – it’s not that the external argument *never* matters. *The comedian killed the audience* has the sense intended above in (15e) (to “wow” the audience, as the gloss in the literature puts it) in part because that’s the thing that can be done with audiences. But, grimly, another thing that can be done with audiences is to cause them to become dead. And so *kill* can have the same meaning it does in (15a= *kill a cockroach*). This meaning might even become more prominent if the external argument were the *killer*, or if you found out that *the comedian* referred to a mass murderer. But what’s going on here is that the internal argument *the audience* makes a few different senses available (the *wow*-sense and the *cause-to-become-dead*-sense), and the external argument disambiguates. But this doesn’t undercut the asymmetry because the external argument’s ability to affect the meaning is still limited. There’s no external argument that you can give to *throw a fit* to get it to have the meaning of *throw* in *throw a baseball*.

³⁶ In the Davidsonian proposal as in the others, adjuncts would be further predicates of events (like WITH(e, a knife)) conjoined with the saturated verbal predicate.

arguments enter into the semantic representation of sentences. Kratzer (1996) proposes that this happens by means of a special composition operation called “Event Identification”.

Using (16d) as a starting point, and abstracting away from tense and aspect, Event Identification takes $\lambda e.\text{STABBING}(e,c)$ and $\lambda e.\text{AGENT}(e,b)$ and conjoins them as $\lambda e.[\text{AGENT}(e,b) \ \& \ \text{STABBING}(e,c)]$.

This is part of what makes Event Identification special. Notice that it takes two λ -bound expressions, and conjoins them while λ -abstracting over the conjoined expression. Simply conjoining $\lambda e.\text{STABBING}(e,c)$ and $\lambda e.\text{AGENT}(e,b)$ would not ensure that the event variables denoted the same eventuality. Event Identification takes the two predicates of events and identifies the two event variables in the computation.³⁷

This means that Event Identification allows for a variadic function operator of a very specific type, and one different from what Recanati had in mind. It takes a semantic unit (the verb and its internal argument(s)) as a functor and yields something which can get an additional argument. Moreover, it’s a fairly ubiquitous kind of variadic function operator in that every sentence containing an external argument will have one. To be clear, Event Identification is just a composition principle; it allows two expressions to compose in a specified way. But the fact that external arguments are not part of a verb’s lexical meaning, and that they compose with the verb by means of Event Identification, means that the semantic operation through which external arguments become part of the semantic representation of a sentence have input and output conditions identical in essential respects to Recanati’s variadic function operators.³⁸

The second point to emphasize has already been made implicitly. Exactly what kind of external argument can be added via Event Identification is constrained by the properties of the event described by the lexical verb and its internal arguments. For example, *kill a conversation* can have an external argument that is an agent (as in *The teacher killed the conversation.*) or perhaps an instrument (as in *The awkward joke killed the conversation.*) Another example, provided by Kratzer, says that *own a dog* will not have an AGENT argument, but an POSSESSOR argument.³⁹ Both points—the ubiquity of variadic function operators like Event Identification, and the constraining

³⁷ In more technical terms, the need for a principle like Event Identification can be explained in terms of a type mismatch. In Kratzer’s theory, this composition happens in a dedicated projection called the Voice Phrase. Voice will first take the verb, then compose it with $\lambda x \lambda e.\text{AGENT}(e,x)$, at which point it can get *Brutus* can fill the argument slot for the agent. So, in terms of the type theory, we have an expression of type $\langle e, \langle s, t \rangle \rangle$ composing with an expression of type $\langle s, t \rangle$ and yielding an expression of type $\langle e, \langle s, t \rangle \rangle$:

$$(i) \lambda x \lambda e.[\text{AGENT}(e,x)]_{\langle e, \langle s, t \rangle \rangle}, \lambda e[\text{STABBING}(e,c)]_{\langle s, t \rangle} \rightarrow \lambda x \lambda e[\text{AGENT}(e,x) \ \& \ \text{STABBING}(e,c)]_{\langle e, \langle s, t \rangle \rangle}$$

If the only composition principle were function application, these expressions wouldn’t compose because of a type-mismatch. But the framework Kratzer is supposing here takes function application to be the mode of composition of lexical arguments, and this process involves composition of non-lexical, functional items, which can have a composition principle other than mere function application.

³⁸ I thank a reviewer for urging me to be clearer on this point.

³⁹ There is some variation on the literature on precisely what thematic roles there are and how they are delineated. In her discussion of this point, Kratzer describes states as having a “holder” external argument. This variability doesn’t affect the overall point, which is that the resulting predicate meaning can constrain the kind of argument that can compose with the verb through event identification. Cf. Carlson 1984, Dowty 1991, Rappaport Hovav and Levin 1998.

role of the event e in determining what kind of argument can be added via Event Identification—will play a role in the Coercion View’s explanation of agentive *ought*.

4.2 Modal semantics and coercion

Although, as I’ve argued, external arguments get added to the semantic representation of sentences through a kind of variadic function operator, it’s important to note that we don’t yet have evidence of the kind of operator that can give us the structure of (13). What last section showed is that it is likely that an instance of Event Identification occurs within the expression ϕ .

(13) OUGHT [OP_{AGENT= x} (ϕ)]

Section 4.1 talked about the unique manner in which external arguments, like agent arguments, compose with verbal expressions. But the Coercion View has it that whenever you have an agentive *ought*-sentence, there is an *additional* operator that adds an agent argument *over and above* the initial external arguments of verb. But you can’t just tack on an additional agent argument willy-nilly into a clause—such an operation needs to be motivated. One way of motivating it is to show evidence of additional structure that allows it. The present section argues that agentive *ought*-sentences have this extra structure. The need for the structure owes to the presence of the modal, and it’s the *coercion* operation that is responsible for it. This section will explain how this comes about, using Kratzer’s modal semantics as a touchstone. Along the way, I will explain how modal semantics is integrated into the kind of event semantics framework supposed in Sect. 4.1.

Though she doesn’t discuss *ought* specifically, it is often analyzed according to Kratzer (1977, 1981, 1991, 2012)’s canonical semantics for modals. This makes it much like the \Box of modal logic, as a quantifier over sets of possible worlds.⁴⁰ One of Kratzer’s innovations is to say that modals are interpreted relative to two kinds of *conversational backgrounds*—the modal base f and the ordering source g , which jointly restrict the domain of quantification. We can schematize the logical form of an *ought*-sentences as follows.

(17) OUGHT (D) ^{$f(w), g(w)$} ϕ

D is the quantificational domain (also called the restrictor), which is determined through context’s selecting a modal base f and ordering source g . f is an accessibility relation on the world of evaluation w . Of the worlds delivered by f , g then selects the highest ranked worlds among those. D is the set of the best ranked worlds (according to g) selected by f . The resulting truth conditions are as follows: (17) is true in w just in case the best g -ranked f -worlds in D are ϕ -worlds.

For deontic modals, the modal base picks out a set of propositions characterizing the circumstances in w relevant to evaluating the modal, and the ordering source picks

⁴⁰ Caveat: I ignore entirely here the fact that *ought* is often said to be “weaker” than other necessity modals, like *must*. This observation was made by Horn (1972). To the extent that this weakness owes to a semantic property of *ought*, it is sometimes said that the domain of quantification is more restricted than for strong necessity modals like *must*. For more discussion, cf. Ninan 2005, von Stechow and Iatridou 2008. This is an interesting and perplexing difference between *ought/should* and *must*, but does not have a bearing on the present discussion, so I put it aside.

out a set of propositions which determine the relevant ideals, laws, or priorities in force in w . Evaluating (18) below, for example, would work as follows.

(18) Jay ought to give up smoking.

The modal base picks out a set of propositions characterizing the relevant circumstances; *that Jay is a smoker, that smoking is expensive, that smoking causes cancer*, etc. Since propositions are themselves sets of worlds, intersecting each of these sets collects the worlds characterizing the conjunction of these propositions in a single set; the set of worlds consistent with the contextually relevant circumstances. The ordering source picks out a set of propositions which characterize the relevant priorities; *that Jay remain healthy, that Jay doesn't spend money frivolously*, etc, and it then ranks the worlds in the domain based on their meeting those priorities. Given these priorities, and holding fast these circumstances, the giving-up-smoking worlds are ranked higher by the ordering source than the continuing-to-smoke worlds. (18) is then true just in case every such giving-up-smoking world is ranked higher than a continuing-to-smoke world. More succinctly: the best worlds in D are Jay-giving-up-smoking worlds.

The account just described follows Kratzer 1981, 1991 in abstracting away from tense and aspect, but there are thorny questions about the interaction of modals with these other clausal elements. It's long been pointed out, for example, that different modals sit at different positions in the clause, and that this results in different scope behavior for these modals.⁴¹ At the same time, it's thought to be an advantage of Kratzerian modal semantics that it provides a *uniform* account of different kinds of modals, deriving different modal "flavors" from a common kernel of meaning through contextually determined values of the modal parameters f and g . The work of Valentine Hacquard addresses the question of how to maintain a uniform, Kratzerian account of modals in light of the height difference exhibited by different modals.⁴² This can be done by minimally re-configuring the conversational backgrounds. Notice in (17) how both f and g were functions from worlds to sets of worlds. Hacquard proposes to keep Kratzer's overall framework, but now f and g are functions from *events* to sets of worlds. In addition to further integrating modal semantics with Davidsonian event-semantics, this allows her to explain the height difference for different modals. Whether a modal base is circumstantial or epistemic depends on the kinds of events the modal base takes as an argument.

Since root modals sit "low" in the clause, closer to the verb, the modal base can take the Verb Phrase's event variable as an argument, whereas on Kratzer's account it took the world of evaluation as argument.⁴³ The truth conditions for (18) remain

⁴¹ Epistemic modals sit "high" in the representation, scoping above tense and aspect, and root modals like the deontics we are concerned with here scope low, below aspect. Cf. Jackendoff 1972, Cinque 1999, Drubig 2001, Stowell 2004, etc.

⁴² Cf. especially her 2006 dissertation and subsequent work, notably Hacquard, 2010.

⁴³ Epistemic modals sit "high" enough in the clause that the VP event variable is closed (already existentially bound, on most analyses) and is no longer "available". The modals instead take the utterance or illocutionary event as an argument. This allows Hacquard to give a principled explanation for why the height difference correlates with the difference in modal bases. Root modals are keyed to VP events and when f takes such an event as an argument, it yields a set of circumstances. Epistemic modals are keyed to events that have content—like illocutionary events or attitude events—and when f takes such an event as an argument it yields an information state. So, it's the kind of argument that determines the value of the function; it's not that the modals themselves require different types of parameters.

more or less the same, except for the subtly different way D is calculated, thanks to the re-configured modal base parameter reflected in (19), which replaces (17). (20) presents a side-by-side comparison of how the circumstantial modal base picks out a set of worlds on Kratzer's and on Hacquard's proposals.

(19) OUGHT (D)^{f(e),g(e)}ϕ

- (20) a. $\bigcap f_{circ}(w) = \{w' \mid w' \text{ is compatible with the contextually relevant circumstances in } w\}$
 b. $\bigcap f_{circ}(e) = \{w' \mid w' \text{ is compatible with the contextually relevant circumstances of } e\}$

But the schema in (19) is still incomplete. According to Hacquard's analysis, root modals scope below tense and aspect, so a more complete representation of the truth conditions for *ought* should include this. Let's fill out (19) with a standard semantics for both of these. Aspect turns on the perfective/ imperfective contrast, and concerns the way the grammar presents or packages the temporal viewpoint of the event described by the sentence.⁴⁴ Let's use the semantics for imperfective suggested by Kratzer (1996), according to which imperfective is characterized as follows: $\lambda t.\lambda e.[t \subseteq \tau(e)]$.⁴⁵ For present tense, we'll say that the truth conditions are defined if $t = t_u$ – that is, if t is identified with the time of utterance. So, for sentences of the form \lceil OUGHT $\phi \rceil$, the schema in (21) replaces (19).

(21) Defined if $t = t_u$

If defined, $=\exists e [t \subseteq \tau(e) \ \& \ \text{OUGHT (D)}^{f(e),g(e)}\phi(e)]$

So far so good. And, not yet any sign of a need for additional structure, much less the kind of structure that would provide evidence for the Coercion View of agentive *ought*.

However, two problems come into relief at this point. Let's call them the orientation problem (OrProb) and the event description problem (EvDescProb) respectively.⁴⁶ First, the OrProb, which is this. Since e is both the variable of the event description *and* the input to aspect, the schema (21) would give truth conditions that make the ϕ event—the event describing the action the subject is supposed to undertake—contemporaneous with the utterance. In a word, it would predict that the interpretation of (21) is *present* oriented instead of *future* oriented. But *ought* sentences with eventive complements are

⁴⁴ Perfective aspect presents the event “from the outside,” as a completed event, whereas imperfective aspect presents it “from the inside,” as an ongoing and incomplete event. Cf. Comrie 1976 for an introduction to these notions and an explanation of the spatial metaphor usually invoked to explain them.

⁴⁵ τ , the temporal trace function, is a function from events to their run-times, or the interval throughout which they occur. Kratzer's proposal here encodes the “from the inside” metaphor for imperfective aspect by saying that a time t (which, incidentally, is the input to tense) is included in the runtime of the event. If the tense is present, it amounts to saying that the event is included in the time of utterance. Equivalently, the event is on-going at the utterance time.

⁴⁶ These problems are discussed in Skibra 2020, though there I called the the second of these the “Event Identification Problem”. (I thank a reviewer for insisting I change the name, so as not to cause confusion with the *composition rule* Event Identification discussed in the paper.) Both problems were first pointed out to me in a slightly different context by an anonymous reviewer for a different journal and of a different paper. I thank them, whoever they are. I should also acknowledge that Homer, 2011 also proposes that root modals need an additional event variable, as I will come to argue here, though on slightly different grounds.

future oriented. So, this representation would give the wrong temporal interpretation of this entire class of *ought* sentences.

Now for the EvDescProb. As indicated in (20b), the modal base uses e to calculate the worlds in D . But (20b) requires that the worlds given by f , the modal base, be consistent with the circumstances of e . For *many* instances of *ought* sentences, this just will not be the case. In these cases, the circumstances according to which *ought* is evaluated are what call for ϕ , and they can be incompatible with ϕ , which is precisely why we need to change them. This is not a problem in standard Kratzerian semantics for modals—Kratzer’s doubly relativizing modals to a modal base and an ordering source effectively avoids this issue. But this very problem pops up in a novel form in Hacquard’s semantics, because of the way deontic modals take the VP’s event variable as an argument.

It is in light of *these* two problems that there arises the need for the kind of additional structure that will motivate the Coercion View. But it is hard to see the need for a coercion operator with the problems stated at this level of granularity. To see this properly, we need to delve into the details a bit more.

4.2.1 The orientation problem

Our stock example in (18= *Jay ought to give up smoking*) can help illustrate both of these problems. Now, as a reminder, the modal base is calculated as in (22) (repeated from (20b), above). Adding tense and aspect into the account gives us the truth-conditions indicated in (23):

- (22) $\bigcap f_{circ}(e) = \{w' \mid w' \text{ is compatible with the contextually relevant circumstances of } e\}$
- (23) $\llbracket \text{Jay ought to give up smoking} \rrbracket^{w,f,g}$ is defined if $t = t_u$
 If defined, $= \exists e [t \subseteq \tau(e) \ \& \ \text{OUGHT } (D)^{f(e),g(e)} \text{ **Jay-give-up-smoking}(e)**]$

We need to have a more fine-grained representation of the meaning of *ought* than we get in (23) for OrProb and EvDescProb to become apparent. This is given in (24), and we can revise the truth-conditions in (23) to (25).⁴⁷

- (24) $\llbracket \text{ought} \rrbracket^{w,f,g} = \lambda P.\lambda f.\lambda g.\lambda w[\forall w' \in \text{BEST}_{g(e)}(\bigcap f(e)):P(e)(w')=1]$
- (25) $\llbracket \text{Jay ought to give up smoking} \rrbracket^{w,f,g}$ is defined if $t = t_u$
 If defined, $= \exists e[t \subseteq \tau(e) \ \& \ \forall w' \in \text{BEST}_{g(e)}(\bigcap f(e)): \text{**Jay-give-up-smoking}(e)(w')=1]**$

Recall that (18) is future oriented. But according to the truth conditions in (25), the time of utterance is contemporaneous with the giving-up-smoking events populating the best worlds. This is so since e is a giving-up-smoking event in all the best worlds, and in those worlds their run-times overlap the utterance time. So our initial schema does not capture the future orientation of (18)—worse still, it would represent the orientation as *present*. This is decidedly *not* the right interpretation of *Jay ought to give up smoking*. If we were to offer this as advice to Jay, and after reflecting on it,

⁴⁷ Remember the characterization of the meaning of *ought* from earlier: the modal takes a modal base, which gives us the worlds compatible with event e , and an ordering source, which ranks these worlds. The truth conditions then “say” that the best of those worlds are worlds where Jay gives up smoking.

he subsequently carried out a resolution to quit, he'd be dutifully complying with this advice, not failing to. The fact that the truth conditions in (25) fail to capture this is the heart of OrProb.

4.2.2 The event description problem

EvDescProb is a bit more subtle. In (25), the event variable is existentially bound outside the scope of the modal (as it needs to be, because this variable is also an input to the imperfective operator). But it is also the event variable contributed by the verb of the prejacent. So the event denoted by e occurs *both* in the actual world and in the deontically ideal worlds. What event is this? In particular, what actual-world event is an ideal-worldly event of giving up smoking?

This raises difficult questions about how to identify an event across different worlds.⁴⁸ Luckily we don't need to answer them fully to be able to compute the truth conditions for this sentence and see that there is a problem. Hacquard gives us a plausible principle to serve as a constraint on what may count as the same event across worlds, the Preservation of Event Description principle (PED).

(26) *Preservation of event description (PED)* for all worlds w_1, w_2 , if e occurs in w_1 and in w_2 , and e is a P -event in w_1 , then *ceteris paribus*, e is a P -event in w_2 as well.

According to PED whatever may count as the same event across different worlds, if it's a giving-up-smoking event in one world, it's a giving-up-smoking event in any others where it exists. But notice what this does to the evaluation of the modal. On Kratzer's original semantics, we were able to say something plausible about how a person's circumstances related to the content of the prejacent; one holds the circumstances of their world w fixed (at a particular time, say), and says of the best worlds consistent with these circumstances, that they are Jay-quitting-smoking worlds. So there is a straightforward sense in which Jay's smoking can be thought to give rise to the obligation. If Jay's smoking weren't part of the relevant circumstances, the worlds in D would be different, and the best of those worlds needn't be ones where Jay quits.

Such a connection between a person's circumstances and the content of the modal cannot be maintained in (25). This is due to the way the modal base is calculated—as a function of the Verb Phrase's event variable. Let's walk through the explanation why. By assumption, Jay's circumstances are as we described them at the beginning of this section. Naturally, we'd want the fact that Jay smokes to be one of the propositions included in the modal domain D . The VP event, e , is a *Jay-giving-up-smoking* event in all of the deontically ideal worlds. But, via PED, it is also a *Jay-giving-up-smoking* event in the actual world! This is quite odd, because our assumption was that Jay's circumstances are characterized in part by the fact that he smokes. Worse still, if actual-world- e were somehow an event of Jay-giving-up-smoking, then $\bigcap f(e)$ could not contain worlds where Jay smokes *now* (since Jay's circumstances cannot both be

⁴⁸ Cf. Hegarty 2010 for relevant discussion.

characterized by his smoking and his giving-up-smoking at the same time). But that's unacceptable. After all, it's *because* Jay smokes that he ought to quit.⁴⁹

Neither of these problems are unique to (25); they both generalize. The OrProb generalizes to *all* future oriented *ought*-sentences. The EvDescProb generalizes to a subset of these where the event description provided by the embedded verb conflicts with a feature of the contextually relevant circumstances. Either of them alone show that the representation given by (25) is unacceptable as truth conditions for (18).⁵⁰

If the meaning of *ought* is as described above, then what results are truth conditions that don't match the natural interpretation of the resulting sentence, for a significant class of *ought* sentences (and modals more generally). Part of the problem is that, on the semantics just described, we need the modal base to project from an event variable located around the VP-level of the clause, but we run into trouble when (a) the verb is eventive and (b) the event variable that is the input to the modal is also the one to which the descriptive material of the verb applies. But once the problem is put in just this way, a fairly simple fix for both problems suggests itself—inserting an additional event variable in the logical form at the VP-level for the modal to project from can give truth conditions which accord with the natural interpretation. *This* is the source of the coercion operator that the Coercion View gets its name from. The location of the proposed coercion operator in relation to the other elements of the clause is given in (27). This effectively gives us a structure like (28), which, one might add, is starting to look a lot like structure we are motivating, given in (13).

(27) [_{TP} PRES [_{AspP} IMPF [_{ModP} OUGHT [_{vP} OP [_{vP} $\phi_{eventive}$]]]]]

(28) OUGHT [OP (ϕ)]

We still need to give the truth conditions for OP and the *ought* sentence that contains it. As suggested earlier, OP needs to add an additional event variable to the truth conditions, and relate this to the embedded verb's event variable in a way that makes clear that former event precedes the latter. Since the former—the newly added event variable—is the input to imperfective aspect and affects the interpretation of tense, this will give us future orientation, solving the OrProb. And since the modal base function

⁴⁹ This problem doesn't arise for Kratzer's account because a *world* can be both a smoking world and a giving-up-smoking world. But an *event* can't be a giving-up-smoking event and still be included in a contemporaneous set of circumstances characterized by smoking. A reviewer raises a good question here: the PED is formulated as a *ceteris paribus* principle. So why doesn't the *ceteris paribus* aspect of the principle kick in to allow that the counterpart of *e* at the actual world is *not* one where Jay gives up smoking? The answer recapitulates a familiar point about *ceteris paribus* conditions; unless we can specify the ways in which things' not being equal is material to the outcome of the principle, the principle will collapse into triviality. For Hacquard, the *ceteris paribus* conditions are linguistic; certain environments will *signal* that these conditions obtain. She makes the case that they are signaled in French by the presence Conditionnel morphology. I have no opinion on whether this works for French, but this kind of morphology does not occur in the examples we are considering in English—nothing signals that *ceteris paribus* conditions obtain. If the only reason to think that they did obtain is that we need them to so that the example works, then the PED would be toothless. Thanks to a reviewer for raising this issue.

⁵⁰ There are fairly easy fixes for the Orientation Problem. In fact, in her dissertation, Hacquard gives a logical form for these problematic sentences (root modal sentences with eventive complements) that gets around the Orientation Problem. Cf. also Kratzer 2010, Matthewson 2012. But the Event Description Problem is a bigger problem, as long as root modals are keyed to the VP event, as they are in Hacquard's system. Skibra (2020) argues that even though they may evade the Orientation Problem, Hacquard, Kratzer, and Matthewson's proposals are still subject to the Event Description Problem.

projects from the new event variable, the modal base will no longer be constrained by the PED, solving the EvDescProb. For now, let’s call this relation R , remembering that R should minimally have a meaning according to which it precedes the more deeply embedded event.

- (29) a. $\llbracket \text{OP} \rrbracket = \lambda e_2. \lambda P. \exists e_1 [P(e_1) \ \& \ R(e_2, e_1)]$
- b. Defined if $t = t_u$
 If defined, $= \exists e_2 [t \subseteq \tau(e_2) \ \text{OUGHT} \ (D)^{f(e_2), g(e_2)} : \exists e_1 (\phi(e_1) \ \& \ R(e_2, e_1))]$
- c. $\llbracket \text{OUGHT} \ [\text{OP} \ (\phi)] \rrbracket^{w, f, g}$ is defined if $t = t_u$
 If defined, $= \exists e_2 [t \subseteq \tau(e_2) \ \& \ \forall w' \in \text{BEST}_{g(e_2)} (\bigcap f(e_2)) : \exists e_1 [\phi(e_1) \ \& \ R(e_2, e_1)] (w')=1]$

In fact, we don’t need to cast around too long for what R might be. The claim being made is that OP is a coercion operation that allows us to capture temporal regularities triggered by aspectual properties of the verb. The work of Dowty (1979) suggests an inventory of abstract predicates in the verbal domain that help model such regularities; for example, CAUSE, BECOME, DO, etc. Such predicates are already in use in semantics to model causative verbs among other phenomena, and they can be put to use in explaining the kind of aspectual coercion that underlies inchoative readings of verbs. So, I suggest that we avail ourselves of something like Dowty’s CAUSE predicate to use in place of R . Adopting this suggestion means we can revise (29) along the lines in (30).

- (30) a. $\llbracket \text{OP} \rrbracket = \lambda e_2. \lambda P. \exists e_1 [P(e_1) \ \& \ \text{CAUSE}(e_2, e_1)]$
- b. Defined if $t = t_u$
 If defined, $= \exists e_2 [t \subseteq \tau(e_2) \ \text{OUGHT} \ (D)^{f(e_2), g(e_2)} : \exists e_1 (\phi(e_1) \ \& \ \text{CAUSE}(e_2, e_1))]$
- c. $\llbracket \text{OUGHT} \ [\text{OP} \ (\phi)] \rrbracket^{w, f, g}$ is defined if $t = t_u$
 If defined, $= \exists e_2 [t \subseteq \tau(e_2) \ \& \ \forall w' \in \text{BEST}_{g(e_2)} (\bigcap f(e_2)) : \exists e_1 [\phi(e_1) \ \& \ \text{CAUSE}(e_2, e_1)] (w')=1]$

To recap the upshot of this section: integrating the semantics of modals with tense and aspect requires of eventive prejacent that they include a kind of coercion operator, which I’ve labeled OP. Importantly, one of the things OP does is to insert another event variable into the semantic composition. This is important for avoiding the OrProb and EvDescProb, and provides an important bit of structure which we will use in our explanation of agentive *ought*.

Before moving on, I’d like to acknowledge an important point. The OrProb and the EvDescProb are both conceptual problems having to do with the relation between the way the modal base is calculated and the modal event (the event described by the prejacent). So, not only does this problem generalize to other eventive *ought*-sentences, but to other eventive modal sentences more generally. The coercion operation I’m describing is a general phenomenon concerning modals, so you’d expect to find it in other root modals. One may worry about this proliferation of coercion operations, but having pointed out the need for it, we can actually start to find evidence of the additional event variable it predicts pretty easily. Cf. (31), for example.

- (31) a. [*Context*: Deadlines have been moved up a week and coursework is now due sooner.]

Yesterday, Pierre still had to turn in his paper next week. Today, he has to turn his paper in tomorrow.

- b. [Context: Yesterday, John’s cupboards were bare, and he needed to go grocery shopping for his party tonight. Before he gets the chance, his friends went grocery shopping for him.]

Yesterday, John had to go to the store today. Today, he no longer has to.

There is a mismatch between the temporal adverbs in each of the sentences in (31); *Yesterday/ next week*; *today/ tomorrow*; *yesterday/ today*. Yet, this mismatch doesn’t result in incoherence or semantic anomaly. It is clear that the second adverb in each pair modifies the event described by the verb in the prejacent. The first adverb modifies the time at which the putative necessity is to hold. This kind of modification is a standard diagnostic for underlying event variables. Cf. Parsons, 1990. The fact that this time can be modified independently of the time of the prejacent is good evidence that an “additional” event variable is present in (root) modals besides *ought* and *should*.^{51, 52}

4.3 Putting the pieces together

Sections 4.1 and 4.2 put us in a position to properly motivate (13). Now, let’s put the ingredients together.

First, an eventive verb embedding under *ought* triggers OP via the posited coercion mechanism.

$$(32) \llbracket \text{OUGHT} [\text{OP} (\phi)] \rrbracket^{w.f.g} \text{ is defined if } t = t_u$$

If defined, =

$$\exists e_2 [t \subseteq \tau(e_2) \ \& \ \forall w' \in \text{BEST}_{g(e_2)} (\bigcap f(e_2)): \exists e_1 [\phi(e_1) \ \& \ \text{CAUSE}(e_2, e_1)] (w')=1]$$

In virtue of OP having applied, the logical form of the sentence now has an additional event variable. This “new” event variable is the structure that can host an additional application of Event Identification, so an additional AGENT argument can compose with the clause in (32), giving us (33).

$$(33) \llbracket \text{OUGHT} [\text{OP} (\phi)] \rrbracket^{w.f.g} \text{ is defined if } t = t_u$$

If defined, = $\exists e_2 [t \subseteq \tau(e_2) \ \& \ \forall w' \in \text{BEST}_{g(e_2)} (\bigcap f(e_2)): \text{AGENT}(e_2, x) \ \& \ \exists e_1 [\phi(e_1) \ \& \ \text{CAUSE}(e_2, e_1)] (w')=1]$

And (13) is just (33), at the relevant level of abstraction.

$$(13) \text{ OUGHT} [\text{OP}_{\text{AGENT}=x} (\phi)]$$

‘OP_{AGENT=x}’ is simply the combination of the coercion operation described above with an AGENT argument conjoining with the new event variable. (It remains to be

⁵¹ I used the semi-modal *have to* in these examples because it inflects for tense, and so it is easier to come up with examples that show this mismatch. Modals like *may* and *must* do not inflect for tense, so cooking up examples that exhibit the mismatch is more difficult, but still possible. Since the coercion operation is an essential part of my account of agentive *ought*, a natural question to ask is whether this means that all these modals admit of agentive readings. I’ll address this question later, but I should point out now that the coercion operation is a necessary, but not sufficient condition for the agentive reading. I thank a reviewer for recommending that I clarify the points made in this paragraph and footnote.

⁵² Homer (2011) specifically provides evidence of this type to show that there is an additional event variable in such root modal constructions, though his examples are from French.

seen what x should refer to in $\text{AGENT}(e_2, x)$. For now we can think of it simply as an unpronounced variable.)

Thus ends my case that (13) is not only linguistically plausible, but also well motivated. It is not an *ad hoc* maneuver to accommodate a complement-ambiguity strategy for explaining agentive *ought*. Each component responsible for giving us the structure in (13) is motivated independently of questions concerning agentive *ought*. Linguistic plausibility aside, it's a separate question whether (13) gives us a workable and attractive account of agentive *ought*. In the final sections of the paper, I will make the case that it does. Prior to moving on, though, let us work through the agentive interpretation of *Jay ought to give up smoking* as the Coercion View would have it.

- (34) $\llbracket \text{Jay ought to give up smoking} \rrbracket^{w, f, g}$ is defined if $t = t_u$
 If defined, $= \exists e_2 [t \subseteq \tau(e_2) \forall w' \in \text{BEST}_g(e_2) (\bigcap f(e_2)) : \& \text{AGENT}(e_2, x) \& \exists e_1 [\text{Jay-give-up-smoking}(e_1) \& \text{CAUSE}(e_2, e_1)] (w') = 1]$

Here is a quasi-English gloss on these truth conditions, letting x refer to Jay: “There is an event e_2 in which the utterance time is contained, such that the worlds which are the best g -ranked worlds consistent with the circumstances of e_2 are worlds where Jay causes an event e_1 of him giving up smoking.” So, given the circumstances—that he smokes, that smoking causes cancer, etc.—realizing one of the most highly ranked worlds (ranked highly according to the standards that he do what he can to stay healthy, that he avoid spending money frivolously, etc.) is matter of him causing a particular kind of event—namely one where he gives up smoking.

Finally, a remark about how we got to this point. The Coercion View depends in part on the claim that root modals sit low in the clause, above the verb but below tense. In turn, the proposal that the modal base and ordering source take events as arguments is meant to accommodate this syntactic fact. That modals do sit low is an empirical claim about English syntax, but it's also one that syntacticians find cross-linguistically robust. (Cf. (Cinque, 1999) and subsequent literature.) This is significant, since it means the Coercion View's explanation of agentive *ought* is available not just to English but to any languages where root modals sit in this “low” position, and the proposal I offer is plausibly a general one about agentive *modals*, not particular to English.⁵³

5 Interpreting the Coercion View

The previous section showed us how the structure that gives rise to agentive *ought* is derived. There are clearly some moving parts to the Coercion View. The goal of the final sections of the paper is to advance the claim that the Coercion View gives us a workable and attractive account of agentive *ought*. The present section will contribute to this aim by showing that these moving parts allow the Coercion View to track some common intuitive judgments about agentive *ought*.

⁵³ Thanks to a reviewer for suggesting I discuss this point more explicitly.

5.1 Coercion without agentivity?

On the Coercion View, OP applies obligatorily to *ought* sentences with eventive prejacents. The additional AGENT argument is optional. It's the combination of the two operations that yield agentive *ought*. So, it's possible for OP to apply without the second application of Event Identification. This is the right result, because we can have non-agentive, future-oriented *ought* sentences, like (2b) and (2e) from our initial suite of examples.

It also means that it's at least a grammatical possibility that all agentive *ought* sentences have a non-agentive interpretation as well. This is also the correct prediction, I think, and one that most competing views also attempt to accommodate, usually through some form of ambiguity. For example, AIP would explain the different interpretations in terms of a different distribution of a verb's theta-roles, which as I argued in Sect. 3 is implausible. Other views, like the *ought*-ambiguity view of Schroeder (2011), would say that the two different interpretations are due to two different *oughts*. The Coercion View claims that whether the *ought* sentence is agentive or not depends on whether a second AGENT argument is present or not. This is a *kind* of ambiguity, but it's a grammatical ambiguity, not a lexical one having to do with the embedded verb or with *ought*.

5.2 What triggers the additional agent argument?

By now it should be clear what triggers OP; an eventive prejacents composing with a root modal. (More precisely: OP is triggered obligatorily by an eventive prejacents. With a stative prejacents, it is optional, but in this case OP would result in future orientation.) But to what do we owe the *second* application of Event Identification? Recall one of the points of Sect. 4.1: what kind of external argument composes via Event Identification is constrained by the properties of the event described by the lexical verb and its internal arguments. That was why *buy a car* can compose with an AGENT argument. The event so described is one that has an agent. *Own a car*, on the other hand, does not get an AGENT argument but a POSSESSOR argument (or a HOLDER argument, according to Kratzer). This is because *own a car* does not describe an event that can count an agent amongst the participants of the event, but a state that has a possessor.

So, assuming OP has applied, the structure exists for Event Identification to apply. When does it? It applies when the underlying event is one that calls for it; when the event is one that has an agent. To see what this means, let's walk through an example—in particular, (10a) (repeated as (35)) which was said earlier to have a prominent non-agentive reading and a remote agentive reading.

(35) Luckless Larry ought to win the lottery.

Since *win the lottery* is eventive, OP would be triggered, yielding the following truth conditions.

(36) $\llbracket (35)_{non-agentive} \rrbracket^{w,f,g}$ is defined if $t = t_u$
 If defined, $= \exists e_2 [t \subseteq \tau(e_2) \ \& \ \forall w' \in \text{BEST}_{g(e_2)} (\bigcap f(e_2))]:$
 $\exists e_1 [\text{LL-win-the-lottery}'(e_1) \ \& \ \text{CAUSE}(e_2, e_1)](w')=1]$

(36) says, roughly, “There is an event e_2 in which the utterance time is contained, such that the worlds which are the best g -ranked worlds consistent with the circumstances of e_2 are worlds where e_2 causes an event e_1 of Larry winning the lottery.” e_2 is the state of Larry’s circumstances. The modal base is a function from e_2 to the set worlds compatible with the contextually relevant circumstances—namely, that Larry is in financial distress, that he’s a a deserving guy, that the lottery being played is fair, and that the Mega Millions or Powerball Lottery is being played next week, etc. In order for e_2 to cause an event of Larry’s winning of the lottery, a number of things need to happen. Larry needs to buy a ticket, decide on the appropriate numbers, etc.

The causal chain that will lead to his winning is one where the circumstances Larry is in cause certain lottery numbers to be drawn from the draw machine. But there is no plausible agent whose machinations bring about this outcome. This is just to say, there’s no agent that brings it about that the Larry wins the lotto. It’s just that a causal chain emerges from the circumstances and results in Larry winning the lottery. Since an eventuality of this kind has no agent responsible for bringing it about, Event Identification would not apply to e_2 to give it an AGENT argument, any more than an event described by the verb *own* would.

Compare with a different context, where (35) does have an agentive interpretation—the admittedly remote one we spoke of in Sect. 3. We can get this interpretation if we suppose that, as part of Larry’s circumstances, he has the knowledge and means by which he can rig the lottery. In this circumstance, the appropriate truth conditions would be (37).

(37) $\llbracket (35)_{agentive} \rrbracket^{w,f,g}$ is defined if $t = t_u$
 If defined, $= \exists e_2 [t \subseteq \tau(e_2) \ \& \ \forall w' \in \text{BEST}_{g(e_2)} (\bigcap f(e_2))]: \text{AGENT}(e_2, x) \ \&$
 $\exists e_1 [\text{LL-win-lottery}'(e_1) \ \& \ \text{CAUSE}(e_2, e_1)](w')=1]$

(37) reflects a circumstance where there *is* a causal chain linking Larry’s circumstances with the outcome of his winning the lottery that involves Larry’s undertaking to make this happen. Indeed, (37) could be evaluated true or false, based on whether g ranks the worlds where Larry successfully rigs the lottery as best or not. But because this e_2 is an eventuality causally linked to the production of another eventuality in a way which invokes someone’s agency, Event Identification can apply to e_2 . This is how we distinguish agentive *ought* sentences from non-agentive (but otherwise future oriented) *ought* sentences. Once again, the Coercion View allows one to explain this difference without an underlying ambiguity in either *ought* or *win the lottery*.

5.3 An unpronounced agent argument?

Since the “additional” agent argument is unpronounced, it’s fair to ask whether positing such unexpressibilia is plausible. I maintain that it is, since unpronounced agent arguments have been posited for various purposes in syntax and in semantics. For example, there is a famous argument that the short passive has an unpronounced agent argument.⁵⁴ The idea is that the short passive in (38a) has an unpronounced agent argument corresponding to the overt *by*-phrase in long passives like (38b). Its presence explains why short passives pattern with long passives with respect to control of “rationale” or “purpose” clauses in (38c) and (38d), and the subject of *collect the insurance* can be the same in the both of them.

- (38) a. The ship was sunk.
 b. The ship was sunk by the owner.
 c. The ship was sunk to collect the insurance.
 d. The ship was sunk by the owner to collect the insurance.

This analysis is not without controversy, but it shows that the positing of implicit agent arguments is not without precedent.⁵⁵

There is a salient difference between the kind of unpronounced agent argument posited for the short passive, and the kind posited by the Coercion View. The Coercion View posits an agent argument over and above the agent argument of the embedded clause. Is there evidence elsewhere that something like this exists? Again, I maintain that there is such evidence. Copley (2008, 2009, 2014) gives a noted account of *futurate* sentences—sentences like (39) that refer to the future in spite of their present-tense morphology.

- (39) a. The Red Sox play the Yankees tomorrow.
 b. The plane leaves at 4:30.
 c. We go to the movies on Thursday.

According to Copley, these sentences have particular felicity conditions; they are only assertible when there is a *plan* at the time of utterance to carry out these acts, and it is common ground who the relevant agent responsible for carrying this plan out is. Moreover, this agent is specified in the logical form of the sentence as an unpronounced AGENT argument (though Copley calls it the “director” argument.) So, it is not implausible to posit an additional agent argument over and above the overt agent argument in the embedded clause. In fact, readers familiar with Copley’s account of futurates will recognize that the Coercion View owes a lot to this work.

5.4 Who can be the agent of *ought*?

Aside from saying it is unpronounced, I haven’t yet said much about what can be the value of the additional agent argument, or how this is resolved. The simplest

⁵⁴ Cf. Faraci 1974, Roeper 1987, Jones 1991, a well as the discussion in Glanzberg 2009.

⁵⁵ Cf. Williams 2015 for critical discussion of this argument and Grimshaw 1990 for related arguments that the short passive has an unpronounced agent argument.

implementation is that the variable be resolved much like a pronoun. In principle, this would allow a certain permissiveness about who could be the “agent” of the *ought*, since the referent of the additional agent argument is whoever can be resolved to be the value of this variable.

Some philosophers might not like this permissiveness. The Coercion View actually offers some degree of flexibility on this, which is useful in light of the considerable disagreement in the literature about *who* can be the relevant agent of an agentive *ought* sentence. Schroeder (2011, pp. 30–33), for example, argues that it can *only* be the subject of the complement. So, for *Jay ought to give up smoking*, the additional agent *x* can only be Jay.

If Schroeder is right, the Coercion View has a variety of ways of accommodating this putative constraint. The first possibility is syntactic. It could be that the argument moves from the original position as the AGENT of the embedded clause to the higher position as the AGENT of the new projection, leaving a trace in the original position. This would mean that the agent of the *ought* and the sentential subject are always the same, thus constraining the interpretation of agentive *ought* as Schroeder would want.⁵⁶

The second is pragmatic. Schroeder makes a further claim, which is that it is *incoherent* if the subject of agentive *ought* were someone or something other than the subject of the embedded clause.⁵⁷ If it is really *incoherent* to have someone other than the embedded subject as the agent of *ought*, so to speak, then we have to worry little about over-generating interpretations of the unpronounced pronoun with someone other than this agent—speakers will not resolve this pronoun with a value that renders the sentence incoherent. The only coherent interpretation will be the one that duplicates the lower agent argument.⁵⁸

However, I’m suspicious of Schroeder’s incoherence claim, and agree with Bronfman and Dowell (2018) and Chrisman (2016) that the agent of *ought* can be someone other than the sentential subject. Let’s presume that the referent of the variable is determined in a similar fashion to pronoun resolution for deictic pronouns. Consider the following scenario:

- (40) [*Context*: Tom and Bill, now happily married, visit a cardiologist for some health issues Tom has been having. The cardiologist is trying to get Tom to make some lifestyle changes for the sake of his health, but knows that Tom is stubborn and set

⁵⁶ Such an account would warrant additional explanation, and I’m agnostic about whether it would ultimately work. For example, it may violate Chomsky (1993)’s Theta-Criterion, according to which each argument bears only one theta-role.

⁵⁷ Cf. his discussion of Broome’s views at p. 31: “Similarly, ‘Jon ought that Jon gets rich’ is ungrammatical, but I don’t think it is incoherent [...]. 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...”

⁵⁸ There are clear examples where semantic incoherence rules out constructions which are otherwise grammatically acceptable. For example (from Williams 2015, p. 69), adjuncts are syntactically iterable. So a sentence with two adjuncts *tomorrow* and *tonight*, like (i) below, is grammatically well-formed, yet incoherent, so we don’t encounter such sentences in natural conversation.

i. #Mo might cook sausage tomorrow tonight.

If it is really *incoherent* for the additional agent to have a different referent than the referent of the subject of the embedded clause, this is not a problem for the Coercion View. It will be ruled out for similar reasons that we don’t see constructions like (i).

in his ways. He knows that, if anyone, Bill can make sure that Tom actually makes some of these adjustments to his lifestyle. After Tom leaves, the cardiologist turns to Bill and says to him.]
Tom ought to exercise daily.

Given the circumstances and the previous discourse, it seems reasonable to think that this is a genuine case of agentive *ought* where the additional agent is Bill instead of Tom. Of course, this claim rests more on intuition than argument, but the important point is that if this *is* an acceptable interpretation of the agentive *ought*, the Coercion View allows for this. One would merely need to maintain that the agent of *ought* is resolved to be the most salient person the discourse participants could coordinate on. This is the view I favor.

That being said, there might be well-defined constraints on how such a pronoun could be resolved which are worth investigating further. To a first approximation, it would seem that the agent of *ought* can't have the status of "discourse new". That is, the candidate agent would already need to have been referred to in the discourse, or be readily inferable from it. In (40), Bill's discourse status as the addressee makes him sufficiently salient to serve as the value of the free variable. As the discussion shows, the framework provides some flexibility on this issue, depending on one's antecedent theoretical inclination.

5.5 What kind of agent is the agent of *ought*?

We might have a good sense of what it means to be the agent of an event described by a particular verb—what it means to be the agent of *give up smoking* or *buy a sun hat*, for example. But what does it mean to be the agent of an eventuality in the sense indicated by the truth conditions of an agentive *ought* sentence as described by the Coercion View? So far, we've read off the truth conditions in a fairly straight-forward fashion, but it's a fair question to ask what this actually means.

What it takes to be the agent of an event described by *buy a sun hat* involves a person's giving a vendor money in exchange for a particular good, since this is roughly what the verbal predicate 'BUYING' (as in BUYING(*e*)) means. But it's significant that the additional event variable supplied by OP has no lexical material like the 'BUYING' or 'GIVING-UP'. It serves as the argument to functional operators like aspectual and modal operators. The question, then, is what sense of agency is required by an event description like ' λe_2 [OUGHT (D)^{f(e₂),g(e₂)}: $\exists e_1(\phi(e_1))$] [Jay-give-up-smoking(*e*₁) & CAUSE(*e*₂, *e*₁)]' for it to be the case that AGENT(*e*₂, *x*)?

My answer comes by way of two remarks—the first conceptual, the second technical. First, though there are no *lexical* predicates applying to *e*₂, the functional predicates applying to it suffice to give us a sense of the agency required. The eventuality *e*₂ is one whereby a person causes another event—a *Jay-giving-up-smoking-event*—in the deontically best accessible worlds. So the sense in which Jay can be an agent of *e*₂ is just the sense required by his being able to bring about a second event—in this case, an event of his giving up smoking. Whether or not a person can be an agent participant of this kind of event depends in part on what the second event is. One can more plausibly be an agent with respect to bringing about an event of one's giving up smoking than an

event of one's winning the lottery, though there may be remote interpretations where one can be an agent with respect to one's winning the lottery.

Second, the technical point. The fact that e_2 has both the 'AGENT' predicate and the modal base function ' f ' applying to it is significant. According to the definition of f in (20b), $\bigcap f(e_2)$ picks out the set of worlds consistent with circumstances of e_2 . One of those circumstances is AGENT(e_2, x). And according to PED, if a e_2 is such that x is its AGENT in w , x will be its AGENT in all w' wherein e_2 occurs. So, assuming that Jay is x in AGENT(e_2, x), all worlds in $\bigcap f(e_2)$ will be worlds where Jay exercises his agency. They will vary in how he does so. In the best of those, according to the truth conditions of (34) given by the Coercion View, he'll do so in a way that causes a giving-up-smoking event. This is a fairly minimal sense of agency required by the Coercion View, but it corroborates our intuitive understanding of agentive *ought* sentences.

This last point is worth emphasizing. The Coercion View explains how linguistic resources from the modal and the verbal domains can be brought to bear on an explanation of the difference between agentive and non-agentive *ought*. The relevant differences include the hallmarks outlined in Sect. 2. Among these is the intuition that agentive *ought* is sensitive to whether or not a person carries out the behavior non-accidentally ("agency sensitivity"). The truth conditions predicted by the Coercion View give us this, but it's otherwise fairly neutral on what constitutes intentional action. What the Coercion View *is* committed to, is a tighter connection between the circumstances that might call for some action, and the performance of this action, and it articulates this connection more explicitly in modal space. But it's not committed to a more substantive theory of intentional action, invoking agent causation, for example, where intentional action consists in an agent's causing an action.⁵⁹

5.6 Agentivity for other modals?

Since what triggers OP is not specific to *ought* and *should*, but is in fact common to all modals, this allows for the possibility of agentive interpretations of other deontic modals like *may*, *must*, etc. I won't discuss these modals extensively here, but I do indeed wish to endorse the view that, insofar as these modals have agentive interpretations, they obtain these interpretations via the same kind of mechanism as *ought* and *should*. Because the agentive interpretations of such modals share the basic hallmarks of agentive *ought*, like future orientation, agency sensitivity, etc., I take it as an advantage of the view that it can be extended to these other modals.

However, other deontic modals do differ from *ought* and *should* in interesting ways. A reviewer points out, for example, that agentive interpretations of deontic *may* and *must* are preferred over non-agentive interpretations, whereas both agentive and non-agentive interpretations are fairly easy to access for *ought* and *should*. I think I share the reviewer's judgments, though explaining this will have to await another occasion. The judgments here are subtle and may be influenced by other factors. It's well documented, for example, that deontic *must* has a well-attested performative use.⁶⁰ It may be that

⁵⁹ I thank a reviewer for suggesting I make this clear.

⁶⁰ The notion of "performativity" at issue here is tricky to explain. But two of the central data points is that (i) performative *must* can bring about an obligation rather than just report on it, and (ii) utterances like

the performativity of deontic *must* influences our judgments about agentivity.⁶¹ Other factors might also play a role in our judgments about agentivity, which need to be sorted out in an investigation of how the Coercion View applies to such modals.

6 Explaining the desiderata

In this section, I will show both that the Coercion View improves on AIP, and that it meets the explanatory desiderata set out in earlier sections.

6.1 Future orientation

We'll start with the easiest one. The Coercion View readily explains why agentive *ought* sentences are future oriented as a rule. In fact, the reason is fairly easy to see—future oriented *ought* sentences are precisely those where coercion via OP has occurred. The Coercion View maintains that (i) agentive *ought* is the result of the combination of OP and a second application of Event Identification, and (ii) the additional structure provided by OP (the additional event argument) provides the grammatical structure that allows for the additional agent argument to be added to the clause. Given this, it's easy to see why agentive *ought* sentences are future oriented. It's because agentive *ought* sentences are a subset of the modal sentences where OP has been applied. By contrast, AIP has no explanation of the future orientation of agentive *ought* sentences. It's simply not in AIP's purview.

6.2 Agency sensitivity

Agentive *ought* was supposed to be agency sensitive. That is, an account of agentive *ought* should be able to distinguish between cases where the action described by the prejacent was executed through a person's capacities as an agent, and those where the putative subject was a mere participant. Recall the example: if Tom intended to kiss Bill, and wound up doing so, but not through an exercise of his agency (because his forming the intention made him so nervous he tripped and planted his lips on Bill), this doesn't seem to capture what the agentive *ought* is counseling in *Tom ought to kiss Bill*. Such an outcome may be quite OK on non-agentive interpretations of *Tom ought to kiss Bill*, though.⁶²

AIP attempted to accommodate agency sensitivity by stipulating that if the prejacent contained an AGENT argument, then the *ought* sentence would be agentive. But this stipulation is too strong, since there is a possible non-agentive interpretation of *Tom ought to kiss Bill*, and *kiss* would still assign an AGENT theta role in this case. And, as I've argued, appealing to verbal polysemy to avoid this result is mistaken.

You must wash your hands, but you won't are infelicitous with performative *must*. In fact, (i) is sometimes leveraged into an explanation of (ii). Cf. Ninan 2005 for discussion.

⁶¹ The matter is complicated by the fact that performativity and agentivity are sometimes conflated in the linguistics literature.

⁶² But cf. the caveats discussed in fn. 15.

The Coercion View has an easier time accommodating agency sensitivity. According to the Coercion View, agentive *ought* sentences contain an additional agent argument, and the truth conditions are such that this agent causes the event described by the prejacent. So, the worlds ranked as deontically best would be those where Tom, as agent, caused the Tom-kissing-Bill event. The accidental-kissing worlds would not be among these on the agentive interpretation. They may be included in the set of deontically best worlds on the non-agentive interpretation, which lacks the additional agent argument. So the Coercion View is indeed agency sensitive.

6.3 Event role asymmetry

Event role asymmetry was the observation that *Tom ought to dance with Bill* may be true while *Bill ought to dance with Tom* false, even though the *dance with* relation is symmetrical. AIP does a pretty good job with argument asymmetry, since *dance* has only one AGENT argument. The problem for AIP is that it can't accommodate non-agentive interpretations of these sentences without appealing to the problematic polysemy claim. Since I've addressed this issue at length, I won't belabor this point further.

Argument asymmetry is easy to explain on the Coercion View, because of the presence of the additional agent argument in agentive interpretations of *ought*. Let's say an utterance of *Tom ought to dance with Bill* has an agentive interpretation. On the Coercion View, this is because there is an additional agent specified in the logical form of the sentence. Let's say Tom causes the Tom-dancing-with-Bill event in the deontically best worlds; these are not worlds where Bill causes the same event. Assuming agentive interpretations of *Tom ought to dance with Bill* and *Bill ought to dance with Tom* where the overt subject of each is also the "additional" agent argument, the truth conditions of these sentences will not be equivalent.

6.4 Over- and undergeneration

There is a bit of overlap between the agency sensitivity, argument asymmetry, and over-/under-generation worries Schroeder and others pointed out for the AIP. This is because AIP's way of dealing with agency sensitivity and argument asymmetry winds up failing to track agentive and non-agentive interpretations of *ought* precisely. But there are other issues lurking in the background, so I'll treat these objections separately. The over- and undergeneration problems do not arise on the Coercion View. We'll return to the *Luckless Larry* sentences to illustrate.

- (41) a. Luckless Larry ought to win the lottery.
 b. It ought to be the case that Luckless Larry wins the lottery.

Schroeder's undergeneration objection was that there was a remote agentive interpretation of (41a) which AIP couldn't explain because *win* doesn't take an AGENT argument (and so no agentive interpretation of *ought* would be triggered). AIP would therefore undergenerate agentive interpretations of *ought*. On the Coercion View, the source of the agentive interpretation of *ought* is not whether *win* has an AGENT argument, but

whether the OP operator and Event Identification have applied. And so the Coercion View can posit this (remote) agentive interpretation *without* positing an ambiguity in *win*, as the AIP was forced to do in an attempt to circumvent the undergeneration objection.

The overgeneration complaint was that AIP will always predict an agentive interpretation if the embedded verb has an AGENT argument. Again, to avoid this overgeneration, AIP is forced to posit the kind of verbal polysemy we've criticized at length. There's no overgeneration of this type on the Coercion View. There's a further worry, though. Schroeder says that (41b) *never* has an agentive reading. If, based on AIP's polysemy claim, we say that (41a) does, then so should (41b), since there's no difference in the thematic roles between the two. Overgeneration looms again for AIP. But nothing in the Coercion View rules out an agentive interpretation of (41b). Does this mean the Coercion View is therefore guilty of overgeneration?

I don't think so. There is some dispute in the literature about whether Schroeder is right to say that (41b) can never have an agentive interpretation.⁶³ Chrisman (2016) and Bronfman and Dowell (2018) argue that it can, given the right circumstances. I'm inclined to agree with them, so I'm less inclined to view this as genuine overgeneration on the part of the Coercion View. Still, it seems clear that when the *ought* sentence is of the form *It ought to be that...* ("it-cleft" sentences), the agentive interpretation is the less prominent one. It would be nice to have a story about why this is.

I don't have a theory of how *ought* functions in *it*-clefts here, but some remarks may suffice to show why the agentive interpretation is less apparent in (41b). Finlay and Snedegar (2014) give some helpful remarks in this direction. They claim that *It ought to be that...* sentences are wordier ways of expressing the corresponding *X ought to...* sentences. So, if the latter have a prominent agentive interpretation, then an application of the Gricean maxim of manner will lead one to disprefer an agentive interpretation for the *it*-cleft sentence.

I think there is more that can be said than this, in particular because *It ought to be that...* sentences are a particular *kind* of *it*-cleft—namely, what Prince (1978) calls "informative-presupposition *it*-clefts". These are known to have a different discourse function from contrastive-stress *it*-clefts since they don't project the same focus structure as the second. (Compare *It ought to be that Luckless Larry wins the lottery* with *It ought to be Luckless Larry that/ who wins the lottery*.) The nature of this discourse function can perhaps explain why such *it*-cleft *ought* sentences rarely get an agentive interpretation. But I leave such investigation to future research.

There's a further issue lurking in the background that is worth addressing. Chisholm (1964) famously argued that sentences of the form α *ought to* ϕ are equivalent to *It ought to be that* α ϕ s—the so-called Meinong-Chisholm Reduction. One way to deny the Meinong-Chisholm Reduction is to say that *It ought to be that* α ϕ s is never agentive, but α *ought to* ϕ is at least sometimes agentive, and so the proposed equivalence fails. This seems at least partially what motivates Schroeder's claim that *It ought to be that* α ϕ s is never agentive. Chisholm uses the equivalence to claim there is nothing distinctive about agentive *ought*.

⁶³ This was noted in fn. 3.

But the Coercion View isn't committed to the Meinong-Chisholm reduction either. The Coercion View claims that α ought to ϕ has both agentive and non-agentive interpretations. Insofar as *It ought to be that α ϕ s* can (perhaps in exceptional circumstances) receive an agentive interpretation means that instances of the former can be equivalent to specific instances of the latter. But that's not what Chisholm was concerned to say. The reason why Chisholm posited the Meinong-Chisholm Reduction was because he thought that every instance of an apparent agentive obligation was equivalent to an impersonal obligation. This is a fairly strong claim, and one that the Coercion View is not committed to. In fact, while Chisholm was saying that sentences that seemed agentive are at bottom equivalent to non-agentive paraphrases, what the Coercion View suggests is that these sentence forms that tend to have non-agentive interpretations can, on occasion, have agentive interpretations. Indeed, if the Coercion View is *right* to say that *It ought to be that α ϕ s* can have an agentive interpretation, then this suggests that alleged contrasts between sentences of the form *It ought to be that α ϕ s* and α ought to ϕ shouldn't be used to underwrite claims about the reduction of agentive *ought* to impersonal *ought*.

In conclusion, the Coercion View is not prone to the undergeneration and overgeneration worries the same way that AIP was. Insofar as it countenances agentive interpretations of *It ought to be that α ϕ s*, one may worry that it is still prone to overgeneration. However, pragmatic factors may explain why agentive interpretations tend to be remote for the *it*-cleft sentences. And in any event, if sentences of the form *It ought to be that α ϕ s* do have agentive interpretations, this fact can't underwrite the kind of reductive ambition animating the Meinong-Chisholm Reduction.

7 Conclusion

I've argued for several claims in this paper. First, that there is a viable complement ambiguity strategy for explaining agentive *ought*, the Coercion View. Second, that the Coercion View is superior to AIP. And third, that the Coercion View is not only linguistically viable, but well-motivated. In making the case, I've also argued that AIP had the right idea in supposing agentive *ought* related to the way agency was encoded in language, though this insight needed to be implemented differently. Furthermore, the Coercion View integrates well with existing theories of modals, and shows Kratzer's theory of modals to be accommodating of agentive *ought*. At bottom, the arguments show that there is a way to defend the view that agentive *ought* is built out of several elements—an idea proposed in some work on deontic logic.

I have argued for the plausibility of the Coercion View largely on linguistic grounds, since these are the grounds on which opponents have shown AIP to falter. This being said, I haven't said terribly much about the logical and meta-ethical consequences of the Coercion View. Since the Coercion View offers an account of agentive *ought* to be taken as seriously as any of its standing competitors, I will mostly leave this task to future work, but for a few closing remarks gesturing at some meta-ethical significance of the Coercion View.

What animates the analysis given in the paper is the following impression. The ease with which we can produce paradigmatic examples of agentive *ought*, contrast

them with examples of non-agentive *ought*, and communicate this fairly easily by means of modal expressions, suggests that language provides an efficient mechanism for encoding the concept of obligation associated with agentive *ought*. Yet, though language may be efficient at encoding and expressing the concept, this doesn't mean that it's altogether transparent exactly how it deploys its resources to do so. I've tried to provide an account of how it does, taking seriously the idea that it does so with the resources at its disposal from the language of modality and from the expression of agency in the verbal domain.

The resulting analysis I provided, the Coercion View, is fairly ecumenical on meta-ethical matters—in particular, on some of those mentioned in the introduction of the paper. For example, a fault line in the on agentive *ought* concerns whether it operates on propositions or whether it relates agents to actions. Several arguments in the literature suggest that whether *ought* or its complement is ambiguous has important consequences for this question. *Ought*-ambiguity views have long been urged to accommodate the idea that agentive *ought* relates agents to actions. The Coercion View falls squarely in the complement ambiguity camp, and its endorsement of the Kratzerian framework might lead one to think that it endorses a line of thought according to which agentive *ought* operates on propositions. In fact, I think more work is needed on this topic to be able to give a decisive answer. Because it seems to me that the Coercion View can actually make good sense of the idea that agentive *ought* relates agents to actions, even as it takes on board a fairly standard modal semantics.⁶⁴ But spelling this out will have to await another occasion.

Acknowledgements This paper is a revised version of the final chapter of my dissertation *Modals and their Arguments*, so it benefited from the comments and support of my committee members; Fabrizio Cariani, Michael Glanzberg, Gregory Ward, Alexis Wellwood, and Justin Khoo. I'm particularly grateful to Fabrizio, Michael, and Justin for numerous discussions and comments on multiple drafts at various stages of this project. I also benefited from conversations with audience members at the second Context Communication and Cognition conference at the University of Warsaw, the 47th meeting of the Society for Exact Philosophy at York University, the 3rd AMPRA conference at SUNY Buffalo, and the Early Lunch Philosophy workgroup at the University of Konstanz. I've had the good fortune to discuss this material with Peter Baumann, David Rey, Ryan Doran, Giulia Felappi, Madgalena Kaufmann, Junhyo Lee, Ethan Myers, and Tadeusz Ciecierski (sometimes at the above conferences). Finally, I thank the anonymous reviewers of this article for particularly careful, detailed, and engaging feedback which improved the paper. Some of the work on this paper was completed while I received support from a grant from the University of Konstanz for the Research Initiative "Universal Preferences for Natural Concepts", which I gratefully acknowledge.

Funding This work was partially funded by a grant from the University of Konstanz for the Research Initiative 'Universal Preferences for Natural Concepts'. Open Access funding enabled and organized by Projekt DEAL.

Declarations

Conflict of interest The author reports no potential conflicts of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included

⁶⁴ So I'm very sympathetic to the remarks in the conclusion of Chrisman 2012.

in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Belnap, N., & Perloff, M. (1988). Seeing to it that. *Theoria*, 54(3), 175–199. <https://doi.org/10.1111/j.1755-2567.1988.tb00717.x>
- Belnap, N. D., Perloff, M., & Xu, M. (2001). *Facing the future*. Oxford University.
- Bresnan, J. (1982). The passive in lexical theory. In J. Bresnan (Ed.), *The mental representation of grammatical relations* (pp. 3–86). MIT Press.
- Bronfman, A., & Dowell, J.L. (2018). The language of “ought” and reasons. In D. Star (Ed.), *Oxford handbook of reasons and normativity*. Oxford University Press. <https://doi.org/10.1093/oxfordhb/9780199657889.013.5>
- Broome, J. (2012). Williams on ought. In U. Heuer & G. Lang (Eds.), *Luck, value, and commitment*. Oxford University Press.
- Broome, J. (2013). *Rationality through reasoning*. Wiley.
- Carlson, G. N. (1984). Thematic roles and their role in semantic interpretation. *Linguistics*, 22(3), 259–280. <https://doi.org/10.1515/ling.1984.22.3.259>
- Castaneda, H. (1975). *Thinking and doing*. D. Reidel Publishing Company.
- Chisholm, R. M. (1964). The ethics of requirement. *American Philosophical Quarterly*, 1(2), 147–153.
- Chomsky, N. (1993). *Lectures on government and binding*. de Gruyter.
- Chrisman, M. (2012). “Ought” and control. *Australasian Journal of Philosophy*, 90(3), 433–451. <https://doi.org/10.1080/00048402.2011.611151>
- Chrisman, M. (2016). *The meaning of ‘ought’: Beyond descriptivism and expressivism in metaethics*. Oxford University Press.
- Cinque, G. (1999). *Adverbs and functional heads*. Oxford University Press.
- Comrie, B. (1976). *Aspect*. Cambridge University Press.
- Condoravdi, C. (2002). Temporal interpretation of modals. In D. Beaver, S. Kaufmann, B. Clark, & L. Casillas (Eds.), *The construction of meaning* (pp. 59–88). CSLI Publications.
- Copley, B. (2008). The plan’s the thing. *Linguistic Inquiry*, 39(2), 261–274. <https://doi.org/10.1162/ling.2008.39.2.261>
- Copley, B. (2009). *The semantics of the future*. Routledge.
- Copley, B. (2014). Causal chains for futurates. In M. Kissine, P. de Brabanter, & S. Sharifzadeh (Eds.), *Future times, future tenses* (pp. 72–86). Oxford University Press.
- Davidson, D. (1967). The logical form of action sentences. In N. Rescher (Ed.), *The logic of decision and action*. University of Pittsburgh Press.
- De Swart, H. (1998). Aspect shift and coercion. *Natural Language & Linguistic Theory*, 16(2), 347–385. <https://doi.org/10.1023/A:1005916004600>
- Dowty, D. (1979). *Word meaning and montague grammar*. Springer.
- Dowty, D. (1991). Thematic proto-roles and argument selection. *Language*, 67(3), 547–619. <https://doi.org/10.2307/415037>
- Drubig, H. B. (2001). On the syntactic form of epistemic modality. Unpublished manuscript: University of Tübingen. Retrieved from <https://www.lingexp.uni-tuebingen.de/sfb441/b2/papers/DrubigModality.pdf>
- Faraci, R. A. (1974). Aspects of the grammar of infinitives and for-phrases. PhD thesis, Massachusetts Institute of Technology.
- Feldman, F. (1986). *Doing the best we can*. D. Reidel Publishing Company.
- Finlay, S., & Snedegar, J. (2014). One Ought too many. *Philosophy and Phenomenological Research*, 89(1), 102–124. <https://doi.org/10.1111/j.1933-1592.2012.00646.x>
- von Fintel, K., Iatridou, S. (2008). How to say ought in foreign. In J. Guéron, J. Lecarme (Eds.), *Time and modality, studies in natural language and linguistic theory* (pp. 115–141). Springer. https://doi.org/10.1007/978-1-4020-8354-9_6

- Geach, P. T. (1982). Whatever happened to deontic logic? *Philosophia*, 11(1–2), 1–12. https://doi.org/10.1007/978-94-011-3352-4_3
- Glanzberg, M. (2009). Not all contextual parameters are alike. Unpublished manuscript: Rutgers University.
- Grice, P. (2001). *Aspects of reason*. Clarendon Press.
- Grimshaw, J. (1990). *Argument structure*. MIT Press.
- Hacquard, V. (2006). Aspects of modality. PhD thesis, Massachusetts Institute of Technology.
- Hacquard, V. (2010). On the event relativity of modal auxiliaries. *Natural Language Semantics*, 18, 79–114. <https://doi.org/10.1007/s11050-010-9056-4>
- Hall, A. (2008). Free enrichment or hidden indexicals? *Mind and Language*, 23(4), 426–456. <https://doi.org/10.1111/j.1468-0017.2008.00350.x>
- Harman, G. (1973). Review of “the Significance of Sense”. *Philosophical Review*, 82(2), 235.
- Hegarty, M. (2010). Two types of belief report. *The Baltic International Yearbook of Cognition, Logic and Communication*. <https://doi.org/10.4148/biyclc.v6i0.1572>
- Higginbotham, J. (1985). On semantics. *Linguistic Inquiry*, 16, 547–593.
- Hilpinen, R. (1973). On the semantics of personal directives. *Ajatus*, 35, 140–157.
- Hilpinen, R. (1974). On the semantics of personal directives. In C. H. Heidrich (Ed.), *Semantics and communication* (pp. 162–179). North-Holland Publishing Company.
- Homer, V. (2011). French modals and perfective. *West Coast Conference on Formal Linguistics (WCCFL)*, 28, 106–114.
- Horn, L. (1972). On the Semantic Properties of Logical Operators in English. PhD thesis, University of California, Los Angeles.
- Horty, J. F. (2001). *Agency and deontic logic*. Oxford University Press.
- Horty, J. F., & Belnap, N. (1995). The deliberative stit. *Journal of Philosophical Logic*, 24(6), 583–644. <https://doi.org/10.1007/BF01306968>
- Humberstone, I. L. (1991). Two kinds of agent-relativity. *Philosophical Quarterly*, 41(163), 144–166. <https://doi.org/10.2307/2219590>
- Jackendoff, R. (1972). *Semantic interpretation in generative grammar*. MIT Press.
- Jackendoff, R. (1997). *The architecture of the language faculty*. MIT Press.
- Jones, C. (1991). *Purpose clauses. Studies in linguistics and philosophy* (Vol. 47). Springer.
- Kaplan, D. (1977). Demonstratives. In J. Almog, J. Perry, & H. Wettstein (Eds.), *Themes from Kaplan* (pp. 481–563). Oxford University Press.
- Klimczyk, J. (2017). ‘Ought’, agents and ambiguity that matters. *Studia Semiotyczne*, 31(2), 113–138. <https://doi.org/10.26333/sts.xxxi2.06>
- Koontz-Garboden, A. (2007). Aspectual coercion and the typology of change of state predicates. *Journal of Linguistics*, 43, 115–152. <https://doi.org/10.1017/S0022226706004464>
- Kratzer, A. (1977). What ‘must’ and ‘can’ must and can mean. *Linguistics and Philosophy*, 1(1), 337–355. <https://doi.org/10.1007/BF00353453>
- Kratzer, A. (1981). The notional category of modality. In H.-J. Eikmeyer & H. Rieser (Eds.), *Words, worlds, and contexts* (pp. 38–74). De Gruyter.
- Kratzer, A. (1991). Modality. In A. von Stechow & D. Wunderlich (Eds.), *Semantik* (pp. 639–650). De Gruyter.
- Kratzer, A. (1996). Severing the external argument from its verb. In L. Z. Johan Rooryck (Ed.), *Phrase structure and the lexicon* (pp. 109–137). Springer. https://doi.org/10.1007/978-94-015-8617-7_5
- Kratzer, A. (2010). What “can” can mean. Unpublished manuscript: University of Massachusetts.
- Kratzer, A. (2012). *Modals and conditionals. Oxford studies in theoretical linguistics*. Oxford University Press.
- Lee, J. (2021). ‘Ought’ and intensionality. *Synthese*, 199(1–2), 4621–4643. <https://doi.org/10.1007/s11229-020-02994-5>
- Marantz, A. (1984). On the nature of grammatical relations. MIT Press.
- Matthewson, L. (2012). On the (non)-future orientation of modals. *Sinn und Bedeutung (SuB)*, 16(2), 431–446.
- Moens, M., & Steedman, M. (1988). Temporal ontology and temporal reference. *Computational Linguistics*, 14(2), 15–28.
- Ninan, D. (2005). Two puzzles about deontic necessity. In J. Gajewski, V. Hacquard, B. Nickel, & S. Yalcin (Eds.), *New work on modality* (Vol. 51, pp. 149–178). MIT Working Papers in Linguistics.
- Parsons, T. (1990). *Events in the Semantics of English. Number 19 in Current Studies in Linguistics*. MIT Press.

- Perloff, M. (1991). Stit and the language of agency. *Synthese*, 86(3), 379–408. <https://doi.org/10.1007/BF00485267>
- Perry, J. (1986). Thought without representation. *Proceedings of the Aristotelian Society*, 137, 137–152. <https://doi.org/10.1093/aristoteliansupp/60.1.137>
- Pietroski, P. M. (1998). Actions, adjuncts, and agency. *Mind*, 107(425), 73–111. <https://doi.org/10.1093/mind/107.425.73>
- Portner, P. (2009). *Modality*. Oxford University Press.
- Prichard, H. A. (1912). Does moral philosophy rest on a mistake? *Mind*, 21(81), 21–37. <https://doi.org/10.1093/mind/XXI.81.21>
- Prince, E. F. (1978). A comparison of wh-clefts and it-clefts in discourse. *Language*, 54(4), 883–906. <https://doi.org/10.2307/413238>
- Pustejovsky, J. (1998). *The generative lexicon*. MIT Press.
- Rappaport Hovav, M., & Levin, B. (1998). Building verb meanings. In M. Butt, & W. Geuder (Eds.), *The projection of arguments* (pp. 97–134). CSLI Publications.
- Recanati, F. (2002). Unarticulated constituents. *Linguistics and Philosophy*, 25(3), 299–345. <https://doi.org/10.1023/A:1015267930510>
- Recanati, F. (2004). *Literal Meaning*. Cambridge University Press.
- Recanati, F. (2010). Pragmatic enrichment. In D. Fara & G. Russell (Eds.), *The Routledge companion to philosophy of language* (pp. 67–78). Routledge.
- Roeper, T. (1987). Implicit arguments and the head-complement relation. *Linguistic Inquiry*, 18, 267–310.
- Schroeder, M. (2011). Ought, agents, and actions. *Philosophical Review*, 120(1), 1–41.
- Sidgwick, H. (1874). *The methods of ethics*. Thoemmes Press.
- Skibra, D. (2020). Future orientation on an event-relative semantics for modals. *North East Linguistic Society (NELS)*, 49(3), 149–162.
- Stanley, J. (2000). Context and logical form. *Linguistics and Philosophy*, 23(4), 391–434. <https://doi.org/10.1023/A:1005599312747>
- Stanley, J., & Szabó, Z. G. (2000). On quantifier domain restriction. *Mind and Language*, 15(2–3), 219–261. <https://doi.org/10.1111/1468-0017.00130>
- Stojnić, U., & Altshuler, D. (2021). Formal properties of “now” revisited. *Semantics and Pragmatics*, 14(3), 1–44. <https://doi.org/10.3765/sp.14.3>
- Stowell, T. (2004). Tense and modals. In J. Guéron & J. Lecarme (Eds.), *Syntax of tense*. MIT Press.
- von Wright, G. H. (1951). Deontic logic. *Mind*, 60(237), 1–15. <https://doi.org/10.1093/mind/LX.237.1>
- Wedgwood, R. (2006). The meaning of ‘ought’. In R. Shafer-Landau (Ed.), *Oxford studies in metaethics* (Vol. 1, pp. 127–160). Clarendon Press.
- Williams, A. (2015). *Arguments in syntax and semantics*. Cambridge University Press.
- Williams, B. (1981a). *Moral luck*. Cambridge University Press.
- Williams, E. (1981b). Argument structure and morphology. *Linguistic Review* 1, 81–114. <https://doi.org/10.1515/tlr.1981.1.1.81>
- Zucchi, S. (1998). Aspect shift. In S. Rothstein (Ed.), *Events and grammar* (pp. 349–370). Springer. https://doi.org/10.1007/978-94-011-3969-4_15

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.