



The Oxford Handbook of Contemporary Philosophy of Language

Ernie Lepore (ed.), Una Stojnić (ed.)

<https://doi.org/10.1093/oxfordhb/9780192856852.001.0001>

Published online: 23 January 2025 **Published in print:** 02 January 2025

Online ISBN:

9780191947445

Print ISBN: 9780192856852

Search in this book

CHAPTER

20 Deontic Modal Expressions

J. L. Dowell

<https://doi.org/10.1093/oxfordhb/9780192856852.013.7> Pages 476–502

Published: 23 January 2025

Abstract

Over the past fifteen years, linguists and philosophers of language have re-examined the canonical, Kratzerian semantics for modal expressions, with special attention paid to their epistemic and deontic uses. This chapter is an overview of the literature on deontic modal expressions. Section 1 provides an overview of the canonical semantics, noting some of its main advantages. Section 2 introduces a set of desiderata that have achieved the status of fixed points in the debates about whether the canonical semantics is correct. Section 3 provides an opinionated overview of some of the main extant rivals to the canonical semantics, including Cariani, Kaufmann, and Kaufmann's complex contextualism; Yalcin's expressivism; Willer's dynamic semantics; and Starr's dynamic expressivism. Section 4 provides an assessment of each of the views discussed in terms of the desiderata introduced in section 2. Section 5 is an overview of remaining issues that require more attention in the literature.

Keywords: [deontic modal expressions](#), [contextualism vs expressivism](#), [deliberative vs evaluative ought](#), [information sensitivity](#), [the miners problem](#), [the Frege–Geach problem](#), [deontic logic](#), [Ross's Paradox](#), [Chisholm's Paradox](#), [weak vs strong necessity](#)

Subject: [Philosophy of Language](#), [Philosophy](#)

Series: [Oxford Handbooks](#)

Collection: [Oxford Handbooks Online](#)

Modal expressions, such as 'necessarily' and 'possibly', are those that qualify the truth or acceptability of a proposition in some way. Such expressions come in a variety of 'flavors': alethic, teleological, epistemic or evidential, and deontic. As a first pass, deontic modal expressions (e.g. 'should', 'must', and 'may') are modal expressions characterized by their distinctive evaluative uses, uses that evaluate an embedded proposition against some standard of ideality. To get a sense for what is distinctive about such uses, consider two contrasting uses of a single modal sentence:

(1) Sobel should be in his office.

It's Tuesday afternoon and a colleague is looking for Sobel. Usually, he's in his office on Tuesday afternoons, but you haven't checked today. You might utter (1) as a way of communicating (roughly) that, given your information, it's most likely he's in his office. Suppose you and your colleague then knock on his door, only to discover that he's not there. But these are his office hours! You might then utter (1) as a way of communicating that his being there is best or required, given his commitments to his students. This second use is evaluative in our sense. It treats his commitments to his students as a standard against which to assess the comparative ideality of possible facts about his location on Tuesday afternoon.

p. 477 Given this feature, deontic modal expressions are a source of interest not only to linguists and philosophers of language but also to philosophers working on ethics, value theory, and the theory of rationality. Here I focus primarily on issues that have received attention in the literature on natural language semantics over the past fifteen years. In section 1, I introduce the canonical semantics for deontic modal expressions, noting some of its main advantages. The dominant view in linguistics has been that modal expressions are quantifiers over sets of possibilities (Kratzer 1991, 2012; Lewis 1975). Necessity modals (like 'must') function as universal quantifiers, while possibility modals (like 'may') function as existential ones. Different modal flavors are determined by different values for the parameters that determine restrictions on a modal's domain of quantification. The values for those parameters are determined as a function of the context of utterance. Whether deontic modals expressions are context sensitive in this way is one issue that has arisen in the recent literature.

In section 2, I turn my attention to desiderata that have achieved the status of fixed points in the debates about whether the canonical semantics is correct. These include the observations that deontic modal sentences have both deliberative and evaluative readings and both information-sensitive and insensitive readings. Adequate resolutions of certain puzzles in deontic logic, such as Chisholm's and Ross's Paradoxes, as well as the Frege–Geach problem for expressivism, have also achieved this status.¹

In section 3, I introduce a representative sample of the rivals to the canonical view on offer—expressivist, dynamic, and contextualist. Contextualist rivals to the canonical view add additional parameters to the formal semantics for deontic modal expressions. In contrast, expressivist and dynamic rivals hold that declarative, deontic modal sentences do not express classical, truth-conditional propositions.² These connections to issues concerning truth-conditional semantics and context sensitivity are both sources of broader interest in this debate about deontic modal expressions.

Section 4 provides an overall assessment of the comparative advantages of these different views along the dimensions introduced in section 2, while section 5 briefly discusses some outstanding issues and notes constraints on plausible responses.

1. The Canonical View

As mentioned above, the dominant view in linguistics has been that modal expressions function as quantifiers over sets of possibilities. Restrictions on the domain of quantification are supplied by the context of utterance. In the case of modal expressions that have a deontic use (e.g., 'ought', 'may', and 'must'), that restriction is two-fold. First, context supplies a value for the *modal base*, f , a function from a world of evaluation, w , to a set of worlds, the *modal background*. On the canonical view, the value for f in the case of deontic modals is *circumstantial*. Roughly, the value for f tells us to look for the circumstances in w that share certain characteristics. The modal background, then, will be the set of worlds alike with respect to those $f(w)$ circumstances obtaining (Kratzer 1991). To illustrate: one common restriction is to those circumstances which fix an agent's action-options in w , for example her abilities, resources, and environment.

p. 478 Second, the possibilities in the modal background are ranked in accordance with the degree to which they conform to some standard of ideality. That standard gets determined by the contextually determined value for the *ordering source*, g . Like f , g is a function that takes a world of evaluation as its argument. Roughly, the value for g tells us to go to w and select the standard that has some characteristic. The canonical view permits a wide variety of values for g . Such a value might select the rules of some club or the laws in some locality in w , for example. It may select some individual's preferences or normative commitments in w . Or it may select the content of morality in w . Those worlds in the modal background that conform to $g(w)$ to the greatest extent are ranked most highly. These most highly ranked worlds make up the domain for the modal (Kratzer 1991).

The canonical semantics for deontic modals enjoys several advantages over extant rivals. First, it shares a formal semantics with modal expressions more broadly. As noted above, modal expressions have a variety of uses: alethic, teleological, and epistemic, as well as deontic. On the canonical view, possibility and necessity modals have a unitary formal semantics, with the different flavors being expressed by different values for f and g . This semantics represents the conventional meaning of modal expressions as comparatively simple, that is simpler than rival views. This fits with evidence concerning the early age at which children typically acquire competence with modal expressions (Matthewson 2016). In addition, the canonical semantics, in contrast to its rivals, enjoys cross-linguistic support (Kratzer 1981).

2. Desiderata

Some observations about the behavior of deontic modal expressions have achieved the status of fixed points in the literature. These provide desiderata any theory of deontic modal expressions must accommodate. Here I mention a few.

2.1 Miners and Information Sensitivity

The need to distinguish between information-sensitive and -insensitive uses of deontic modals has been widely recognized in the recent literature.³ Parfit's miners scenario is the most discussed example illustrating that need. Here's a standard statement:

Ten miners are trapped either in shaft A or in shaft B, but we do not know which. Flood waters threaten to flood the shafts. We have enough sandbags to block one shaft, but not both. If we block one shaft, all the water will go into the other shaft, killing any miners inside it. If we block neither shaft, both shafts will fill halfway with water, and just one miner ...will be killed.

p. 479 A widespread judgment is that (2) expresses the correct answer to the question of what to do in this case: 4

(2) We ought to block neither shaft.

Since we know that in the circumstances, blocking neither is suboptimal, (2) must be correct relative to the information in the scenario. In this sense, (2) is information-sensitive.

Suppose, though, that before we need to arrive at a decision, we discover that the miners are in A. In that case, we are now strongly inclined to accept (3) instead:

(3) We ought to block A.

This change in our judgment illustrates the way in which information-sensitive uses are ‘*seriously information-dependent*’: Changes in an agent’s information can make a difference to which of her options is best. Some have concluded that serious information-dependence must not only be representable within a semantic theory for such expressions but also incorporated into their semantics. Some argue that the canonical semantics is unable to capture information-sensitive readings, as in (2) (Kolodny and MacFarlane 2010; Cariani et al. 2013). We will return to this issue in section 3.

In contrast, some uses are clearly information-insensitive. Suppose we are wondering what would be best in the original miners scenario, given the circumstances. In that case, the following sounds fine:

(4) Either we ought to block A or we ought to block B. We just don’t know which.

What is best in those circumstances is to save all of the miners. (4) expresses this.

A common conclusion from Parfit’s case is that deontic modal claims are uniformly information sensitive (see, e.g. Kolodny and MacFarlane (2010) and Yalcin (2012)). However, the felicity of (4) shows that this is not so. We’ll return to this issue, too, in section 3.

2.2 Evaluative vs Deliberative Uses of Deontic Modals

Recently, much attention has also been paid to the distinction between the so-called ‘deliberative’ and ‘evaluative “ought”s’. Many discussions rely upon Mark Schroeder’s (2011) characterization of the distinction. According to Schroeder, the ‘deliberative “ought”’ is characterized by five hallmarks. Such an ‘ought’

- (i) ‘matters directly for advice’;
- (ii) ‘is the right kind of thing to close deliberation’, that is, it ‘settles the question of what to do’;
- (iii) is the one tied to responsibility, that is, what one can be praised or blamed for complying or failing to comply with;
- (iv) is constrained by what is in ‘one’s power to do’; and
- (v) is ‘more closely connected to’ the notion of obligation (than the ‘evaluative “ought”’)

p. 480 For example, if I’m at a racetrack deciding whether to bet on Blue Blazer or Exploder, and I know that in the past, the former has reliably proven the faster horse, I might conclude:

- (5) I ought to bet on Blue Blazer.⁴⁴

There are important distinctions in the neighborhood of Schroeder's hallmarks that can be sharpened and made less controversial. First, it would be less controversial to characterize the distinction as a distinction in the use of deontic modal sentences than to characterize it in terms of distinct 'oughts'. Second, there are scenarios in which there are distinct and incompatible deontic modal claims, each of which has some of these hallmarks but neither of which has all of them. Consider a case in which what it is advisable to do is not the same as what it is most reasonable for an agent to do, given their information. For example, suppose, again, that I am at the racetrack. My information is as before. But now someone more informed overhears my deliberations. This individual knows that the race has been rigged—Blue Blazer has been drugged—and thus Exploder is most likely to win. In that case, she can say,

(6) That woman should bet on Exploder.

In such a case, what it is advisable for an agent to do ((6)) is distinct from what they would be rationally criticizable for failing to do ((5)). In other words, here, Schroeder's hallmarks (i) and (iii) come apart. Which of these deserves the title of 'the' deliberative use? A more useful taxonomy recognizes these further distinctions. I'll call uses as in (6) 'advisability uses' and reserve 'deliberative uses' for those like (5), which will be characterized primarily in terms of his hallmark (ii).⁵ On this way of marking the distinction, deliberative and evaluative uses of deontic modal sentences are distinguished by the sorts of question they may serve as answers to. Deliberative uses, in contrast to evaluative ones, are full or partial answers to questions of what to do, where such answers are treated as *normative* or *action-guiding*.⁶

Their potential action-guidingness is another much-discussed feature of deontic modal expressions. What is it for a use to be action-guiding in the reserved sense? Three distinct interpretations are available. First, there is the idea that to be action-guiding, the acceptance of a deontic modal claim must be *reason-giving* (see, e.g. Wallace 1998; Svavarsdottir 1999). Second, there is the idea that it is to be *motivating* (Blackburn 1988; Starr 2016) and, third, that it is to be *treated as reason-giving for conversational purposes* (Silk 2016). Glossing over these distinctions, we might say that to be minimally action-guiding, the acceptance of a deontic modal claim must 'point' an agent in the direction of some action. The distinct interpretations may then be treated as precisifications of this general notion.⁷

p. 481 In contrast, evaluative uses merely rank a state of affairs along some measure of ideality. For example, a researcher studying the global distribution of existing food resources might conclude: ↵

(7) There ought to be less famine than there is.⁸⁸

Prima facie, (7) represents states of the world in which there is less famine than there actually is as more ideal than the actual state of affairs. It does not settle a question of what to do or serve as a guide to action. Capturing these two distinctive uses of deontic modals is a further constraint on an adequate semantics and pragmatics.

2.3 Deontic Logic

Another source of recent dispute concerns deontic logic and some of its classic puzzles. Two puzzles that are widely thought to place constraints on a plausible semantics and pragmatics for deontic modals are Ross's Paradox and Chisholm's Paradox. Ross's Paradox⁹ concerns the interaction between deontic disjunctions such as (8) and free choice.

Suppose you have promised to mail an important letter for a friend. In that case, (8) seems felicitous:

(8) You ought mail the letter.

Now consider

(9) You ought to either mail the letter or burn it.

Intuitively, (9) does not clearly express what you ought to do. It suggests that there are two ways to do as you ought, burning the letter being one of them. Any semantics, such as the canonical one, which validates *Inheritance*, however, holds that (8) entails (9):

Inheritance: If φ entails ψ , *ought* φ will entail *ought* ψ .

Accounting for the data in Ross's Paradox and other puzzles¹⁰ that target *Inheritance* are widely thought to require revisions to the canonical semantics (e.g., Cariani 2013; Lassiter 2017; Starr 2016).

p. 482 Chisholm's Paradox targets two further inference rules, *factual detachment* and *deontic detachment*¹¹: \perp

Deontic detachment

- (DD i) Ought φ .
- (DD ii) If φ , then ought ψ .
- (DD iii) Therefore, ought ψ .

Factual detachment

- (FD i) φ .
- (FD ii) If φ , then ought ψ .
- (FD iii) Therefore, ought ψ .

Here's a typical statement of the paradox:

- (CP i) Jones ought to help his neighbor.
- (CP ii) If Jones helps his neighbor, he ought to tell them he's coming.
- (CP iii) If Jones doesn't help his neighbor, he ought not tell them he's coming.
- (CP iv) Jones doesn't help his neighbor.

Most are able to consider a scenario in which (CP i)–(CP iv) strikes them as true. If those judgments are all sensitive to a single point of evaluation, and deontic detachment is valid, then (CP i) and (CP ii) together guarantee the truth of

(CP v) Jones ought to tell his neighbor he's coming

relative to that same point of evaluation. (See, e.g., Willer 2014; Arregui 2010; and Saint Croix and Thomason 2019 for discussion.) Given that Jones does not help him, however, (CP v) strikes many as unacceptable in this scenario. (Call these 'the Chisholm judgments'.) Instead, (CP iii) and (CP iv) together incline them to accept

(CP vi) Jones ought not tell his neighbor he's coming.

Factual detachment licenses the inference from (CP iii) and (CP iv) to (CP vi). Since (CP v) and (CP vi) are incompatible,¹² Chisholm's Paradox appears to pit these two inference rules against each other. Since many find (CP vi) more intuitively compelling than (CP v), a common reaction is to hold onto factual detachment, while giving up deontic detachment. Since the canonical semantics validates the latter but not the former, some take this as evidence that the canonical semantics cannot be correct (Arregui 2010; Lassiter 2017; Cariani 2021).

Whether either of these puzzles poses a genuine challenge to the canonical semantics is an issue we will return to in section 3.

2.4 Compositionality and Frege–Geach

p. 483 *Compositionality* is a near-universally accepted constraint on any semantic theory. The idea is that the meaning of a complex expression should be a function of the meaning of the simpler expressions that make it up, together with the compositional rules. One test of a semantic theory for an expression, then, is its embedding behavior in larger \hookrightarrow constructions. According to expressivists about deontic modal expressions, the use of sentences in which the modal takes widest scope does not express representational states of mind. This thesis is a little surprising, given the declarative form of such sentences and that sentences with declarative form are paradigmatic examples of representational discourse. For example, suppose, in explaining why our friend Alex is happy, I say,

(10) Pat is in town.

(10) represents the world as being a certain way. But suppose I say instead,

(11) Alex may skateboard.

On the face of it, (11) has the same representational function as (10). The difference is that, in the latter case, unlike the former, the world is represented in terms of what it makes permissible. One reason to think this hypothesis is on the right track is that (10) and (11) share embedding behavior. Expressivists, however, do not take this to be decisive evidence against their central thesis. The challenge for expressivists is to explain that behavior. This is the *Frege–Geach problem* for expressivism. Mark Schroeder (2008, 2015) has argued that the most difficult such challenge stems from its alleged inability to explain mixed disjunctions, disjunctions of representational sentences, such as (10), with putatively non-representational ones, such as (11). Suppose, for example, I am uncertain why Alex is happy but have limited the possible explanations down to two. In that case, I might say,

(12) Either Pat is in town or Alex may skateboard.¹³¹³

According to Schroeder, the challenge for the expressivist is to explain, compositionally, what overall state of mind sentences like (12) express. In W. Starr's version of the challenge, they asks what an expressivist could hold is communicated by a sentence like (12). I will return to these challenges to expressivism in discussing expressivist proposals in section 3 and again, in assessing them in section 4.

3. Rivals to the Canonical Semantics

3.1 Complex Contextualism

p. 484 Fabrizio Cariani, Magdalena Kaufmann, and Stefan Kaufmann (2013) have developed what is perhaps the least revisionary update to the canonical semantics. Their semantics is also contextualist. But, in addition to the usual modal base and ordering source, they add a parameter δ representing a decision problem to contexts of utterance.¹⁴ This decision parameter is introduced to capture the information sensitivity of deliberative \sqsubset uses exhibited in (2). The basic idea is that the value for δ partitions worlds in the modal background into a set of cells (subsets of those worlds) where each world w in each cell is alike with respect to which of an agent's action options that agent performs in w . We may then think of each cell as a proposition, the proposition that the agent chooses the particular option, of those open to her, that distinguish that cell from the others. In deliberative uses, the ordering source g then ranks cells instead of worlds.

In addition, values for g are purely qualitative. For example, in the miners case, the ordering source delivers

10 miners save < 9 miners saved < 8 miners saved ... 0 miners saved.

To capture serious information dependence, values for f are epistemic. Putting these together, an agent's options are ranked to the extent that the relevant qualitative priorities are guaranteed, given the contextually relevant information (Cariani et al. 2013: 247–248). This is, in effect, to encode a MaxiMin decision rule¹⁵ into the semantics.¹⁶

Capturing serious information sensitivity is one of their proposal's significant advantages. In the original miners case, for example, the option that guarantees the best worst-case outcome is doing nothing. Thus, the theory correctly predicts the truth of (2) and also of (3) in the case in which we learn that the miners are in A. A second advantage is its ability to capture both deliberative and non-deliberative uses, such as (4). In non-deliberative uses, the value for f may be epistemic or circumstantial. Moreover, the value for g will rank worlds, not options. What is best, given the circumstances, according to the qualitative ordering source, is to save ten lives. Blocking the shaft they are in (either A or B, we don't know which) would guarantee that. On that reading, (4) comes out true.

One difficulty for the view as stated concerns low probability, low ranked outcomes. To see this, consider a version of the miners scenario in which we learn that there is a 99.99 per cent chance the miners are in A. (Indeed, perfect certainty being rare, this version of the scenario is nearer a real case.) Since blocking A does not guarantee the highest-ranked outcome, the Cariani, Kaufmann, and Kaufmann semantics still predict that (2) is true. But, from a normative point of view, it does not seem true. In that case, that we ought to do nothing. Instead, it seems that (3) is true—we ought to block A. This isn't a fatal objection to the overall semantics. However, it does show that the view requires revision.¹⁷

3.2 Pragmatic Expressivism

p. 485 The Cariani, Kaufmann, and Kaufmann semantics is representationalist. ‘Ought φ ’ represents the possible options or world states in which φ is the case as having a certain property, namely as being best relative to a set of parameter values. The remaining rivals to the canonical semantics discussed here are non-representationalist. Like contextualists, Seth Yalcin defends a static semantics for deontic modal expressions. \hookrightarrow A traditional, static semantics treats contents as representational truth-conditions. On a *representationalist* view, the content of deontic modal sentences on an occasion of use represents a way a world can be. Such a semantics is then supplemented with a separate pragmatic theory explaining how the acceptance of utterances with such contents change the context for subsequent discourse. Like traditional views, Yalcin’s semantics relegates an utterance’s update effects to a separate pragmatic theory. But unlike traditional views, Yalcin eschews representational contents. The fundamental idea motivating his view is an expressivist one, namely that the states of mind expressed by the use of deontic modal sentences are not representational but plan-laden. This is a pragmatics-first approach. We begin with the states of mind expressed by the use of such sentences and the conversational effects of their acceptance. We then work backward to identify the semantics they must have in order to play their pragmatic role (Yalcin 2012: 132–135).

To home in on the sort of state of mind expressed by a deontic use of a sentence like (1), Yalcin considers what would make a sentence like (13) true:

(13) The chair believes that Sobel should be in his office.

The plan-laden states of mind which make sentences like (13) true are modeled using a set H of hyperplans. A hyperplan is a view about what to do, given an information state. An information state is modeled as a set s of possible worlds representing a choice situation. A hyperplan h is a function that takes an information state s and delivers a subset of that state, namely those worlds in s with permissible outcomes according to h (Yalcin 2012: 147). Following Gibbard (1990), Yalcin builds from these ingredients a view about what it is for a proposition to be required, forbidden, or permissible:

Requirement: realizing a proposition p is *required* in s just in case for every hyperplan $h \in H$, $h(s) \subseteq p$.

Forbidden: realizing a proposition p is *forbidden* in s just in case for every hyperplan $h \in H$, $h(s) \subseteq \sim p$.

Permission: realizing a proposition p is *permissible* in s just in case for every hyperplan $h \in H$, $h(s) \cap p$ is non-empty.

To be unopinionated about the deontic status of p is for none of these conditions to obtain.

The compositional semantics for ‘ought’ Yalcin proposes to capture Gibbard’s ideas is the following:

$$[[ought\varphi]]^{w,h,s} = 1 \text{ iff } \forall w' \in h(s) : [[\varphi]]^{w',h,s} = 1$$

The addition of an information state to the points of evaluation ‘ought φ ’ is sensitive to is motivated by the observation that deontic modals sometimes exhibit serious information dependence, as observed in connection with the miners case discussed in section 2 (Yalcin 2012: 148, 150).

p. 486 Yalcin’s expressivism is captured in his pragmatic story about how unembedded deontic modal claims update a conversational state. First, the needed conversational states will be plan-laden information states, represented by a pair of an information state s and a set of hyperplans, H . Conversational states represent joint states of mind, roughly what is jointly presupposed about that the world is like and what is jointly

planned.¹⁸ The pragmatic function of a deontic claim is to add a constraint to the hyperplans in H (Yalcin 2012: 150). For example, updating the conversational state $\langle s, H \rangle$ with 'must $\sim p$ ' will result in a state $\langle s, H' \rangle$ such that for every hyperplan $h \in H'$, $h(s) \subseteq \sim p$. We may think of accepting 'must $\sim p$ ', then, as jointly planning to rule out realizing p -worlds.

As noted earlier, a central objection to expressivist proposals for deontic language rests on the claim that the needed semantics cannot be made compositional. Yalcin's proposal suggests, however, that this objection rests on a confusion. Compositionality is a constraint on a semantics, not a pragmatics. There is nothing particularly expressivist about Yalcin's semantics. As he himself notes, it is compatible with non-expressivist views about the function of deontic language (2012: 148). That said, it will be helpful to consider in detail how his view might treat the conversational effect of mixed disjunctions, thought by some to present the most difficult embedding challenge to expressivism (Schroeder 2008 and Starr 2016). I will return to this issue in section 4.¹⁹

3.3 Dynamic Semantics

In contrast to Yalcin, Malte Willer (2014) proposes a dynamic semantics for deontic modal expressions.²⁰ A dynamic semantics, in contrast to a static one, integrates the update effects of an utterance's acceptance into the semantics. Such semantics treat the meaning of an expression as its *context-change potential*, a function from contexts of utterance to contexts of utterance. We may think of an expression's context-change potential as a set of instructions for updating a context of utterance. Replacing truth at a point of evaluation, the central notion for a dynamic semantics is acceptance in a context. Very broadly, a context accepts a sentence p just in case updating with p does not result in an empty context. For a simple illustration, suppose we represent a context of utterance as a set S of information states s . Consider now a context such that for each $s \in S$, $s \subseteq \sim p$. In that case, updating each s with p will result in the empty set. We may say, then, that a set S of information states s in which $\sim p$ is true throughout will not accept p .

Willer represents contexts as a set S of information states s . An information state s is represented as a set of possible worlds. Deontic contexts, in addition, determine a deontic ordering source o . An ordering source is represented as a set of propositions, where propositions are represented as sets of possible worlds (Willer 2014: 6, 9). The ordering source identifies an ideal subset of s for each $s \in S$. It does this by ranking worlds. A world w is strictly more deontically ideal than another w' relative to o ($w <_d w'$) if and only if (i) for every proposition $\varphi \in o_d$, if $w' \in \varphi$, then so is w and (ii) there is some proposition $\varphi \in o_d$ such that $w \in \varphi$ and w' is not. A world w is among those deontically ideal in s (s_d) just in case there is no other world in s strictly more ideal than it (Willer 2014: 9). \hookrightarrow A context $\langle S, o \rangle$ accepts 'must φ ' just in case for each s in S , the deontically best subset of s , s_d , is a subset of φ (Willer 2014: 10).

p. 487

Willer argues that one advantage of his semantics is that it solves Chisholm's Paradox while also validating both factual and deontic detachments. Instead, Willer's semantics is non-monotonic. This non-monotonicity allows the view to capture serious information dependence. This means that while (CP i) and (CP ii) entail (CP v) by deontic detachment, that inference is subsequently defeated by the additional information that Jones does not help. Factual detachment then supports the inference from (CP iii) and (CP iv) to (CP vi) (Willer: 2014: 1–3).

However, a semantics that validates factual detachment is not clearly desirable. To illustrate, consider a version of Forrester's Gentle Murder Paradox (Forrester 1984). We might find all of (GM i–iii) acceptable, thus rejecting (GM iv):

- (GM i) No one should murder anyone.
- (GM ii) Jones will murder his neighbor.
- (GM iii) If Jones murders his neighbor, he should murder him painlessly.
- (GM iv) Jones should murder his neighbor painlessly.

But by factual detachment, (GM ii and iii) entail (GM iv). While some may find (GM iv) clearly acceptable here, that judgment is far from universal. Intuitively, the blanket prohibition against murder applies even to those determined to do so and even after the fact. To express that, we want available a true reading of (GM v) in this context.

(GM v) Jones should not murder his neighbor—painfully or painlessly.²²¹

It is not clear how to make the needed reading of (GM v) available in Willer’s semantics.

3.4 Dynamic Semantic Expressivism

W. Starr proposes an expressivist, dynamic semantics for deontic modals. On that proposal, “the meaning of an expressive sentence is characterized in terms of how its acceptance changes a preference ordering” (Starr 2016: 361). Central to their motivation for moving from a static to a dynamic semantics are their claims that (a) there is nothing distinctively expressivist or non-representational about Yalcin’s semantics (Starr 2016: 372–373) and (b) putting together Yalcin’s proposal regarding the states of mind expressed using descriptive language, as in (10), with his proposal regarding the states of mind expressed using deontic language, as in (11), results either in a view according to which the use of mixed disjunctions like (12) either do not communicate anything at all or a view on which ‘or’ has merely pragmatic significance (Starr 2016: 372–373). These claims will be assessed in section 4. For now, let’s focus on Starr’s positive proposal.

p. 488 Dynamic meanings, recall, are functions from contexts to contexts. To model deontic language, Starr represents a context as a state S , where S is made up of a set of substates, $\sqcup \{s_i, \succeq_j\}$. Each s represents a *state of information*, a set of worlds compatible with what interlocutors are mutually presupposing. Each \succeq_j is a *preference frame* made up of two relations on preferences, strict and equal preference, which comparatively rank the worlds in s . The substates that make up a context are in competition. Each information state s is competing to be what is presupposed for the purposes of the conversation. And each preference frame \succeq_j is competing to be the preference ordering which motivates interlocutors (Starr 2016: 378).

Starr proposes to capture ‘the expressivist’s guiding slogan’ that ‘must ϕ ’ ‘promotes certain motivational attitudes towards ϕ ’ as follows. First, ‘must ϕ ’ updates each preference frame in each substate so that the worlds in s that are ϕ are strictly preferred to any that are $\sim\phi$. Second, the result is tested for whether it makes ϕ a practical necessity in each substate (roughly, whether it makes all of the strictly preferred worlds in s ϕ -worlds). If so, ‘must ϕ ’ updates the context. Since preference frames represent joint preferences and preferences are motivational, the result is a context in which interlocutors are jointly moved towards realizing ϕ over $\sim\phi$ possibilities (Starr 2016: 381).

Starr’s proposal nicely captures our sense that deontic discourse can be action-guiding. However, as mentioned in section 2, some uses of unembedded deontic modal sentences are clearly not action-guiding. Recall (7):

(7) There ought to be less famine than there is.

(7) is not action-guiding. Yet, the modal is clearly deontic. Perhaps Starr could suggest that, although not motivational, (7) still serves to express a strict preference for less famine than there is over the actual state

of affairs. However, consider another example. Imagine one feminist telling another about the rules of a local all-men's club she has recently learned of. In listing the rules she says,

(14) Members must not bring women to networking events.

Certainly, (14) sounds felicitous in such a scenario. 'Must' here does not have an epistemic reading. It is deontic. It serves a communicative purpose. But its acceptance isn't preference-updating for the interlocutors in this conversation. Feminists strongly disprefer excluding women from networking events! Building a preference-updating function into the semantics for deontic discourse seems to deprive the resulting language of an ability to communicate what is required according to a set of rules, whenever compliant states are dispreferred in a context.²²

4. Assessments

How well do each of these views fare with respect to capturing the desiderata from section 2? Let's consider them in turn.

p. 489 4.1 Miners and Information Sensitivity

It is easy to see how the canonical semantics can capture information-insensitive readings of deontic modal sentences, as in (4):

(4) Either we ought to block A or we ought to block B. But we don't know which.

Here, the ordering source g ranks each world w in the modal background in accordance with how many miners are saved by the action performed in w . The modal background is circumstantial. The circumstances are that the miners are all in a single shaft, either A or B. So the worlds in the modal background are either all worlds in which the miners are all in A or worlds in which they are all in B. Depending upon which of those possible circumstances are actual, either all of the most highly ranked worlds are those in which A is blocked or those in which B is blocked. So, (4) comes out true.

Pace Cariani, Kaufmann, and Kaufmann (2013), the canonical semantics can just as easily capture the information sensitivity of uses like (2) in the miners case:

(2) We ought to do nothing.

It does this by permitting evidence-sensitive values for g . The deliberative question raised in the miners case is: which action can be expected to save the most lives, given our evidence regarding the miners' location? The relevant circumstances are those that fix our evidence regarding that location and regarding the different possible outcomes for each of our choices. They also include the circumstances that fix those choices, such as the limited number of sandbags available. Given those circumstances, and the goal of choosing the action that can be expected to save the most lives, doing nothing, ranks most highly. So, (2) comes out true.²³

The Cariani, Kaufmann, and Kaufmann semantics is also able to capture both readings. On their view, deliberative readings are always information sensitive. As discussed previously, they capture the information-sensitive readings with epistemic values for f , combined with a value for the decision parameter δ that partitions worlds in the modal background into an agent's action options. Evaluative readings may take either epistemic or circumstantial values for f and g ranks worlds, rather than action options. This allows them to capture purely evaluative uses like

(7) There ought to be less famine than there is.

Worlds in the modal background will be like the actual world with respect to facts that determine the number of calories it is possible to sustainably produce and distribute. $G(w)$ ranks each world in the modal background w comparatively, depending upon how well it does at minimizing famine. Since the minimal famine worlds have less famine than the actual world, (7) comes out true.

p. 490 One drawback of this account of the distinction between deliberative and evaluative uses is that it is unclear how it can capture disagreement over which deontic modal claims are information-sensitive. Consider an objective and a subjective consequentialist disagreeing about what morality's primary requirements dictate in the miners scenario. The subjective consequentialist believes that those requirements are information sensitive so that (2) expresses those dictates. The objective consequentialist believes that those requirements are information insensitive so that (4) expresses their dictates. According to Cariani, Kaufmann, and Kaufmann (2013), however, (2) and (4) are compatible since their contents are determined relative to different parameter values. It is puzzling, then, how our consequentialists could use (2) and (4) to express their disagreement over what morality primarily requires. After all, they each acknowledge the truth of both sentences under the readings given by the Cariani, Kaufmann, and Kaufmann view.²⁴

In contrast to their contextualist rivals, the semantics of Yalcin, Starr, and Willer each build in information-sensitivity. This is because what plays the role in each of their views of the ordering source in the canonical view—hyperplans, preference frames, and orderings, respectively—are all operations on an information state. This makes it difficult to see how to capture the needed information-insensitive readings of (7) or of (4) in the miners case. One option would be to give a separate semantics for deontic sentences that take such readings. But positing an ambiguity in deontic modal expressions is empirically less attractive than the unitary semantics for all modal expressions given by the canonical view. And here, too, it would be difficult to see how such a view could adequately represent disagreement over which deontic modal claims are information-sensitive.

4.2 Deliberative vs Evaluative Uses of Deontic Modals

While it is clear how our expressivist views can capture deliberative or action-guiding uses of deontic modal sentences, it is much less clear how they can capture purely evaluative readings in which a proposition is assessed for its comparative ideality relative to a standard that, in context, may be accorded no normative or action-guiding status whatsoever. For example, as we saw in section 3, it is not clear how a dynamic expressivist like Starr can represent felicitous readings of sentences like (14), which, in context, play no action-guiding role. There is a similar puzzle for Yalcin's Gibbard-style expressivism. For Yalcin, when

(14) Members must not bring women to networking events

p. 491 is accepted into a conversational record, it functions to update the hyperplans that are live in context by eliminating any plan h which, relative to s , permits members bringing women to networking events. But, granted that there are planning states of mind of the kind Yalcin's account of deontic language appeals to, they do not seem to be ones \hookrightarrow updated by the acceptance (14) into the conversational record of our feminists in the example above. Intuitively, what (14) communicates in that context is information about the comparative ranking of members bringing women to networking events and their not doing so, given the club's rules. Thus, the use of (14) seems representational, not motivational.

What about Willer's dynamic, non-expressivist view? That view ties what is best according to an ordering source to the set S of information states in the context. For this reason, it's not clear how his semantics is able to represent evaluative readings that express obligations that it is known (relative to S) will not be

fulfilled. For example, worried about his immortal soul, a Catholic father may say to his wife about their son, Sam,

(15) He won't go to confession, but he ought to.

Or, on learning he won't go, Sam's mother can lament,

(16) But he must! (*Pace Ninan 2005*).

On his semantics, neither (15) nor (16) can be accepted by any context in which the possibilities in the modal's prejacent are already ruled out, relative to each $s \in S$. However, Willer introduces a technical fix for these cases. It involves an operation he calls 'downdating', which adds the missing possibilities into each $s \in S$. While this solution may capture our judgments about the felicity of cases like (15) and (16), it is not clear how the availability of this operation interacts with the validation of factual detachment. The intuition that supports factual detachment is the idea that our information can make a difference to what it makes sense to do; that is, it is the intuition that at least some deontic modal claims are seriously information dependent in the sense introduced in section 2.1. The idea is that once we learn that Jones will not help his neighbor, the possibilities in which he does are no longer relevant for our thinking about what would be best for him to do. The introduction of an operation that sometimes requires the consideration of possibilities which are ruled out by our information and that sometimes does not in order to fit with our judgments about the different cases is a bit ad hoc. Better would be to provide a uniform account that explains how the intuitive readings get selected in the different cases.

4.3 Deontic Logic

Both Willer and Starr cite the non-monotonicity of their views as an advantage in resolving Chisholm's Paradox. (See also Cariani 2013.) Similarly, their rejection of Inheritance permits resolutions of puzzles like Ross's Paradox, which target Inheritance (Starr 2016 389–390). If invalidating Inheritance and embracing non-monotonicity were the only way to adequately resolve these puzzles, each of these \downarrow views would enjoy a significant advantage over the canonical view, which both validates Inheritance and is monotonic. However, invalidating Inheritance is not required to resolve Ross's Paradox.²⁵ And it is far from clear that non-monotonicity is desirable. (See Bronfman and Dowell 2018 and von Fintel n.d.) Finally, as we've seen, there are independent reasons to reject factual detachment, which the canonical semantics does not validate. This provides the canonical view with its own solution to Chisholm's Paradox.

4.4 Compositionality and the Frege–Geach Problem

It is straightforward for any representationalist semantics, like both our contextualist views, to explain the embedding behavior of deontic modal sentences. Yalcin (2012) discusses how such sentences embed under attitude ascriptions in his expressivist-friendly semantic framework. How, though, do each of our expressivist views do at explaining the embedding behavior of deontic modal sentences in mixed disjunctions, like

(12) Either Pat is in town or Alex may skateboard?

Capturing this behavior is an advertised advantage of Starr's account. First, 'or' is given its usual dynamic meaning, the union of the updates associated with each disjunct (Starr 2016: 375).

Next, recall that Starr represents a context as a state S , where S is made up of a set of substates, $\langle s_i, \succeq_j \rangle$. Each substate is made up of an information state (i.e. a subset of the set of all possible worlds) and a preference

frame, which compares the worlds in s for their preferability. Let p = Pat is in town and q = Alex skateboards. A context that accepts (12) will be a set S'' , which is the result of updating S with the union of the sets, S' and S'' , associated with the updates of (10) and (11), respectively. S' will be the set of substates that results from removing all of the $\sim p$ -possibilities from each $s \in S$. The resulting set of substates will be ones in which p is true at each world in each information state. S'' will be the set of substates that results from adding a strict preference for q over $\sim q$ to each preference frame in the substates in S . To capture the effect on S of updating with (12), we simply take the state S'' which is the union of S' and S'' . This context will be one in which interlocutors are not fully opinionated about either Pat's location or the permissibility of Alex's skateboarding, as desired.

In section 2, we saw that Starr does not see a solution in this neighborhood available to expressivists like Yalcin, who combine a static semantics with an expressivist pragmatics. What should we make of this claim? Section 3 noted that Frege–Geach challenges to Yalcin's expressivism are misplaced. Compositionality is a constraint on semantic theories, but Yalcin's expressivism is captured by his pragmatics. Still, we might wonder, what is the conversational effect of the acceptance of a sentence like (12)? What joint state of mind does the resulting conversational state represent? This \hookrightarrow is essentially what Starr asks when they ask what a sentence like (12) communicates on Yalcin's view.

First, let's see what Starr's challenge is and then see whether there is a way for a view in the ballpark of Yalcin's to meet it. Starr argues that because the first disjunct does not place a constraint on hyperplans, it does not rule any out. Likewise, the second disjunct does not place any constraint on states of information, so it, too, doesn't rule any out. Since communication occurs by ruling out elements of contexts, they conclude that mixed disjunctions on pragmatic expressivist views like Yalcin's can't communicate anything (Starr 2016: 373). Is that correct?

Yes and no. Mixed disjunctions do place constraints on Yalcin's conversational states but, for reasons having to do with technical features of his view, the constraints are not ones that could clearly communicate anything. To see this, first recall that a context or conversational state on Yalcin's proposal is a pair made up of an information state and a set of hyperplans, $\langle s, H \rangle$. A context that accepts (12) will meet either one of two conditions. Either it will be one in which $s \subseteq p$ or it will be one in which every $h \in H$ is such that $h(s) \cap q$ is non-empty. Thus, no conversational state $\langle s, H \rangle$ which fails to meet one of these conditions will accept (12). The difficulty is that either a context already meets one of these conditions, so that accepting (12) does not communicate anything, or it doesn't, and (12) does not provide any instruction on which of these two conditions to implement—do we update s ? Or do we update H ?

That said, this new difficulty is due to technical features of Yalcin's view that are inessential to the basic ideas behind it. To capture the mixed conversational states we're after, we will need to slightly enrich Yalcin's states. Let us borrow an element of Starr's proposal. For Starr, contexts are made up of a set of substates. Let us now represent a conversational state C in Yalcin's framework as a set of $\langle s_i, H_j \rangle$ pairs. Call each such pair a 'substate'.²⁶ The pragmatic function of a deontic sentence ϕ is now to add a constraint to each of the hyperplans h in H_j , for each substate $\langle s_i, H_j \rangle$. The pragmatic function of a descriptive sentence ϕ is to eliminate worlds from s_i in each substate $\langle s_i, H_j \rangle$ that are incompatible with its content. We may now build the overall conversational update associated with the use of a sentence like (12) out of the updates associated with each disjunct and their disjunction.

Here are the conditions on an updated conversational state that accepts (12) in this revised framework. Letting p = Pat is in town and q = Alex skateboards,

- (i) 'Pat is in town' is accepted by any conversational state C such that each s_i in each substate $\langle s_i, H_j \rangle$ is such that $s_i \subseteq p$;
- (ii) 'Alex may skateboard' is accepted by any conversational state C such that, for each substate $\langle s_i, H_j \rangle$,

each $h \in H_j$ is such that $h(s_i) \cap q$ is non-empty;

(iii) (12) is accepted by any conversational state C that is a union of substates meeting either (i) or (ii).²⁷

p. 494 This makes the update associated with (12) both precise and non-empty. Starr's objection to this sort of solution on behalf of the pragmatic expressivist seems to be that it that would give 'or' a wholly pragmatic meaning. But this is not so. It gives 'or' a pragmatic \hookrightarrow conversational effect that is compatible with its having its usual truth-conditional meaning. The only new element will be that, on Yalcin's semantics, contents are true and false relative to enriched indices, which require a hyperplan and a state of information in addition to a world.²⁸

This suggests that what is really at issue in deciding between Yalcin's and Starr's proposals is not how to build a plausible expressivist account of deontic language but whether to accept a dynamic semantics or a static one with an expressivist pragmatics.²⁹

5. Outstanding Issues

There are several outstanding issues concerning deontic modal expressions that require further work. Here are a few.

5.1 Semantic Neutrality and Normative and Metanormative Disagreement

Several authors have converged on the idea that a semantics for deontic modals should, in some way, be neutral with respect to how options or worlds get ranked. Nate Charlow, for example, argues that, while deontic modals are relative to a decision problem and a choice function, their semantics should be otherwise neutral with respect to which choice functions they are relative (Charlow 2018).³⁰

Since not all deontic modal sentences are deliberative, though, we should not expect them all to be relative to a decision problem or choice function. Consider again

(7) There ought to be less famine than there is

and

(14) Members must not bring women to networking events.

These are not deliberative; they are not answers to questions about what to do. The kind of ordering-source neutrality we need is stronger than Charlow's.

To see what sort of semantic neutrality would be desirable, I suggest we start with the reasons why such a constraint is attractive. Like any semantics, a semantics for deontic modals should figure in plausible explanations of the use-facts we find. These include their ability to state a wide variety of normative and metanormative positions regarding value, morality, and rationality. They also include the ability of advocates of these various positions to disagree and debate their truth.³¹ For this reason, I propose

p. 495

Semantic neutrality:

A semantics for deontic modal expressions should be neutral between which normative and metanormative positions regarding value, morality, and rationality are correct to the extent such neutrality is possible.

By building risk aversion into the semantics for deontic modals, the Cariani, Kaufmann, and Kaufmann (2013) view violates this constraint. As we saw in the example between our subjective and objective consequentialists, any view that builds information sensitivity into its semantics will also violate this constraint. One advantage of the Kratzerian semantics is its catholicism with respect to the types of value for g that can be supplied as a function of contexts of utterance. In the case of our subjective and objective consequentialists, for example, she can say that their dispute is over what we should do in the miners case, given what morality primarily requires. That locates their dispute where it belongs—in the substantive, normative realm rather than questions about natural language semantics.

5.2 How Is Context Sensitivity Resolved?

Many of the considerations raised above suggest that the canonical contextualist semantics remains the odds-on favorite. But if contextualism is correct, how are the needed parameter values supplied as a function of context? In my (2011) and (2013), I offer an initial proposal in a framework that is Kaplanian in spirit. Kaplan suggests that speaker's intentions play an important role in fixing the denotations of demonstratives (Kaplan 1989). Noting that there is a *prima facie* puzzle about how private mental states could play a role in communication, I suggested that the needed role in fixing modal parameter values could only be played by publicly manifestable intentions. These would be communicative intentions that are made manifest by publicly available features of a conversational situation. Such features could be made manifest by prior discourse, as with a question under discussion. Or they could be made manifest by perceptually salient features of a conversational situation together with shared general world knowledge, as when a discussion is taking place between a doctor and a patient in a medical examination room.

In her (2021), Una Stojnić defends a rival view. According to Stojnić, modals receive their parameter values entirely by linguistic convention, where the relevant conventions include discourse coherence relations. Much of Stojnić's evidence for her proposal is compelling. That said, I am currently inclined to think that linguistic conventions cannot fully explain every case, even if they can explain many. Consider, for example, a busy doctor, who walks into an examination room containing a patient with a visible skin ailment, hands him a prescription and says,

(15) You should take this.

p. 496 There does not seem to be a linguistic convention which determines that the modal in (15) is relative to the patient's physical condition or that each of his options o are being ranked in accordance with the degree to which o improves it. These seem to be provided by general world knowledge that doctor and patient share about the point of going to a doctor's office, a doctor's professional role in such situations, and jointly perceptible facts about the patient's physical condition.³²

5.3 Generating Predictions

As we've seen, a significant test of a semantics and pragmatics for deontic modal expressions is its ability to plausibly fit with robust and widespread speaker judgments about cases. While 'fitting with' merely requires that the needed readings are available within a semantic framework, *plausibly* fitting with our judgments requires an independently plausible story about why the needed readings are to be expected, compatible with the semantics, in the target cases. As just noted, some work on providing such a story for contextualists has been done. For non-contextualists, more work remains to be done. In the case of a dynamic semanticist à la Willer, we need an independently plausible explanation of the difference between cases in which the information of the context constrains which modals can be accepted and cases in which it does not. (This is to say, we want an explanation whose plausibility is independent of the mere fact that some operation, like Willer's selective use of downdating, fits with our judgments.) We also need to know how contexts determine ordering sources. In the case of our expressivists, we need some explanation of how non-motivational, merely observational readings, as exhibited by (14), could get selected in a framework that ties the conversational effect of unembedded deontic modal sentences directly to motivation.

5.4 Weak vs Strong Necessity Modals

It is widely recognized that there is some distinction between weak and strong deontic necessity modals, for example between 'ought' and 'must'. But what that distinction is, and whether and how to capture it in the formal semantics or pragmatically, are subjects of ongoing debate. Exhortations like the following are illustrative of the contrast:

(16) You ought to help your elderly neighbor, but you don't have to.

(16) does not express a contradiction. It might be tempting to explain this by a context shift (perhaps induced by 'but', which marks contrast). However, that would not explain why (17) is marked:

(17) You ought to help your elderly neighbor, but it's not the case that you ought to.

p. 497 A number of different proposals have been floated to account for this. Kai von Fintel and Sabine Iatridou (2006) have suggested a secondary ordering source that further restricts the domain of quantification for 'ought'. In (16), for example, the domain for 'must' might be restricted to worlds in which you fulfill your obligations. The domain for 'ought' would be further restricted to that set of worlds which meet some secondary desideratum, for example of meeting your obligations in some way that would do the most good.

Dilip Ninan (2005) and Paul Porter (2007) have suggested that strong necessity modals exhibit a distinctive directive behavior shared with imperatives. Ninan suggests there is a difference between the felicity of (18) and (19):

(18) Sam should go to confession, but he's not going to.

(19) Sam must go to confession, but he's not going to.

According to Ninan (2005: 2), while (18) is fine, (19) is marked. The difference, he argues, is that while 'must' exhibits a practical force akin to an imperative, 'ought' does not.³³ Both Ninan and Portner propose to model this difference by positing a 'to-do' list as an element of conversational scoreboards. The suggestion is that, just as an accepted imperative adds a commitment to an addressee's to-do list, an accepted 'must' claim does so as well. The explanation of the contrast, then, is that 'must φ ' requires that φ is a live possibility in the context—we cannot add the impossible (relative to the common ground) to an agent's to-do list.

However, these judgments about (18) and (19) are far from robust and universal. Moreover, ‘must’ is often used to express obligation, and obligations do not disappear even when we know that they will not be complied with. (19) is felicitous under just such a reading. A second observation Ninan (2005) makes is perhaps more suggestive of the contrast. In the scenario just mentioned, while it is ‘epistemically impossible’ that Sam goes to confession (it is conversationally presupposed that he won’t), he still possesses the *capacity* to do so. Perhaps ‘must’ has a stronger action-guiding role by being more closely tied to ‘can’ than ‘ought’ is. To see why this might be so, consider a revised version of the Chisholm Paradox sketched above. Suppose I have promised to help my friend in another city. To do so at the appointed time, I needed to leave my home five hours ago. I now tell you of my predicament. Consider these two possible replies:

(20) You ought to have kept your promise to your friend.

(21) You must’ve kept your promise to your friend.

(20) sounds fine on a deontic reading. Sentences like (21), though, Ninan notes, seem incapable of a deontic reading. To get its felicitous reading, ‘must’ve’ requires an epistemic interpretation. Perhaps the difference is that keeping my promise is no longer something that I can do and ‘must ϕ ’ requires the ability to realize ϕ .

p. 498 5.5 An Inferential Constraint on Deontic Necessity Modals?

Jessica Rett (2016) offers strong evidence that deontic necessity modals exhibit an interesting behavior shared by their epistemic counterparts. Others have observed that the strong epistemic necessity modal is infelicitous in contexts in which the speaker’s evidence for the prejacent is direct. Suppose, for example, that I am watching a downpour from my office window. It is not (humor aside) felicitous for me to say,

(22) It must be raining.

The standard observation is that epistemic ‘must’ requires support by indirect evidence. Rett notes, though, that some forms of indirect evidence likewise make (22) infelicitous. Suppose a colleague tells me that it is raining. Testimony is a form of indirect evidence. Yet, if I pass this information along to another colleague, it would be infelicitous for me to do so with (22). Rather, I should simply report that it’s raining. Suppose instead, though, that I utter (22) after observing others enter my building with wet raincoats and umbrellas. My evidence here is not merely indirect; it is inferential. In this case, (22) is felicitous. The key observation, she concludes, is not that epistemic ‘must’ merely requires *indirect* evidence but that it requires *inferential* evidence.

Rett suggests that the deontic necessity modals share this behavior. Both ‘ought’ and ‘must’ require inferential support by salient practical reasons.

6. Summing Up

At the moment, the canonical semantics for deontic modal expressions remains the one with strongest empirical support. It is the only one that is part of an overall, unified semantics for modal expressions, making it the simplest. This fits well with what is known about the early age at which children acquire modal vocabulary. Its comparative simplicity also allows it to capture the great variety of readings we find, deliberative and evaluative, information-sensitive and information-insensitive, action-guiding and non-action-guiding. Unlike rivals that build information sensitivity into the formal semantics, the canonical view is easily able to capture the needed information-insensitive reading of (4), as well as represent disputes over which deontic modal sentences are information-sensitive as normative, rather than linguistic, disputes. Unlike rivals that build a parameter for a planning or preference state into the semantics, it is able to capture the needed, non-action-guiding uses of (7) and (14). Moreover, it fits with independently well-motivated solutions to puzzles in deontic logic and faces no special issues concerning compositionality or the conversational effects of using deontic modal sentences. Finally, the canonical view enjoys strong cross-linguistic support.

p. 499 That said, much work remains to be done. For example, the comparative advantages of dynamic versus static semantics needs further exploration. The dynamic views we have considered all reject classical propositions (which have truth-conditions relative to worlds only) as the contents communicated by the use of deontic modal sentences. However, in recent work, Una Stojnić has shown that assigning dynamic meanings to modal sentences is compatible with assigning them classical truth-conditional contents on an occasion of use. Such an approach to deontic modal sentences has promise that requires further exploration (see, e.g. Stojnić 2021).

Further, all of the views we have discussed require further work to become fully prediction-generating. As we have seen, although work has been done (see, in particular, Dowell 2011, 2013; Lewis 2020; Stojnić 2021), further testing of contextualist hypotheses requires more precise accounts of how parameter-values get determined as a function of contexts of utterance. In the case of expressivist and dynamic views, more work needs to be done in connecting the semantics and pragmatics with our judgments about cases. These views require not just revisions that would allow them to capture the readings they currently render unavailable. In addition, they require an account of how the different readings get selected. For example, in the framework of a semantics or a semantics and pragmatics that build in motivational updates, what generates the needed non-motivational uses in (7) and (14)? Before rejecting any of the views we have considered, these issues would need to be further explored.

Notes

1. The Jackson and Pargetter (1986) puzzle involving Professor Procrastinate also has this status. It's standardly cast as a puzzle for the principle, Inheritance, discussed below in connection with Ross's Paradox. For relevant discussions of the Professor Procrastinate puzzle, see Cariani (2013) and Bronfman and Dowell (2018).
2. Classical propositions in the sense reserved here are those that have truth-values relative to a world alone.
3. This distinction roughly maps onto the metaethicist's distinction between so-called 'subjective' and 'objective "ought"s'. See, e.g., Wedgwood (2016).
4. The example is drawn from MacFarlane (2014: 285).
5. What about the additional hallmarks, like (iii)–(v)? Setting aside (iv) for now, I observe that uses that fail to have (ii) may nonetheless exhibit hallmarks (iii) and (v). For example, on certain views about the status of morality, one is obligated to

comply with the most general moral principle and criticizable for failing to do so. A deontic modal sentence expressing that principle would have hallmarks (iii) and (v) without (ii).

6. This distinction will cross-cut Schroeder's own. He argues that deliberative and evaluative 'ought' sentences are marked by distinct syntactic structures. For an assessment of his empirical claims in support of that proposal, see Chrisman (2012) and Bronfman and Dowell (2018).
7. For additional discussions of deontic modal expressions and action-guidingness, see Ninan (2005) and Portner (2007).
8. See Dowell (2013) for further discussion of uses of this type.
9. Ross (1944).
10. Jackson and Pargetter's (1986) Professor Procrastinate case is another puzzle that targets Inheritance.
11. Chisholm (1963).
12. 'Incompatible' here means 'not jointly true at any point of evaluation'.
- p. 500 13. The example is from Starr (2016: 373).
14. For other complex contextualist proposals, see Cariani (2013, 2016).
15. Such a rule identifies the option, among an agent's set of options, that guarantees the most valuable minimum possible outcome as most choiceworthy.
16. For a critical discussion of this sort of proposal, see Carr (2015).
17. In light of considerations of this kind, Charlow (2018) proposes a revised semantics that adds both a decision problem and a choice function to the points of evaluation to which deontic modals are sensitive.
18. More precisely, Yalcin holds that a sentence ϕ is part of an overall conversational state if the state of mind of each participant reflects the update associated with ϕ . When ϕ is an unembedded deontic modal sentence, that update will be reflected in the set H of hyperplans of each participant (2012: 133).
19. For an alternative static, preference-based expressivist proposal, see Silk (2014).
20. For an account of modal expressions that is both dynamic and representational, see Stojnić (2021).
21. For Willer's discussion of this paradox, see his 2014: 19–20.
22. Starr (2016: 381 fn 44) does distinguish between 'descriptive' and 'performative' uses of deontic modals. However, descriptive uses are those for which ϕ is already strictly preferred over $\sim\phi$ throughout the preference frames of the context. (14) is neither descriptive nor performative in Starr's senses.
23. For further discussion and details regarding information-sensitive readings in the canonical framework, see Dowell (2013) and Bronfman and Dowell (2018).
24. One response would be to hold that our consequentialists are engaged in a linguistic dispute and so are not contesting the truth of either (2) or (4) under the Cariani, Kaufmann, and Kaufmann readings. However, that response is a little strained. As those engaged in these debates see it, their dispute is a normative matter rather than a linguistic matter. The hypothesis that their dispute is, nonetheless, linguistic will require positing fairly widespread semantic ignorance in these cases—amongst those we might expect to be experts in the nuances of moral language!
25. See Bronfman and Dowell (2018) for an Inheritance-preserving solution to Ross's Paradox.
26. Here, too, Yalcin might borrow Starr's reasons for representing conversational states as sets of substates, namely the idea that each substate is competing with the others to be the overall joint state of mind represented by the conversational state.
27. To get the constraints on the conversational states that accept a deontic version of Schroeder's (2015) central example, "either my parents lied to me or one must not steal," we replace p with "my parents lied to me" in (i). Letting q = one steals,

we replace in condition (ii) that $h(s_i) \cap q$ is non-empty with $h(s_i) \subseteq \sim q$. As before, the disjunction is accepted by any conversational state made up of substates meeting either (i) or (ii).

28. A further complaint might be that the truth-conditional content for 'or' will guarantee that any conversational state in which one of the disjuncts is accepted is a conversational state in which the disjunction is accepted and that this is counter-intuitive. However, that result will be found in any semantics that assigns 'or' its usual truth-conditional meaning and assigns that meaning an update effect on conversational states. Thus, there is no special problem for the expressivist here.
29. For discussion of the issues related to choosing between a dynamic semantics and a static semantics plus pragmatics, see Lewis (2020).
30. Other advocates of neutrality include Cariani (2014) and Carr (2015).
31. For a nice discussion of the different ways a semantics might represent two interlocutors as disagreeing, see Rieppel (2011).
- p. 501 32. For a nice discussion of how general world knowledge might play a role in resolving context sensitivity, see Lewis (2020).
33. Alex Silk (2016) also suggests that the difference between 'ought' and 'must' is a difference in their practical force.

References

- Arregui, Ana. 2010. "Detaching If-Clauses from Should," *Natural Language Semantics* 18: 241–293.
[Google Scholar](#) [WorldCat](#)
- Blackburn, Simon. 1988. "Attitudes and Contents," *Ethics* 98: 501–517.
[Google Scholar](#) [WorldCat](#)
- Bronfman, Aaron, and Dowell, J. L. 2018. "The Language of 'Ought' and Reasons," in *The Oxford Handbook of Reasons and Normativity*, Daniel Star (ed.), Oxford: Oxford University Press, 85–112.
[Google Scholar](#) [Google Preview](#) [WorldCat](#) [COPAC](#)
- Cariani, Fabrizio. 2013. "'Ought' and Resolution Semantics," *Noûs* 47/3: 534–558.
[Google Scholar](#) [WorldCat](#)
- Cariani, Fabrizio. 2014. "Attitudes, Deontics, and Semantic Neutrality," *Pacific Philosophical Quarterly* 95: 491–511.
[Google Scholar](#) [WorldCat](#)
- Cariani, Fabrizio. 2016. "Deontic Modals and Probabilities: One Theory to Rule Them All?," in *Deontic Modality*, Nate Charlow and Matthew Chrisman (eds), Oxford: Oxford University Press, 11–46.
[Google Scholar](#) [Google Preview](#) [WorldCat](#) [COPAC](#)
- Cariani, Fabrizio. 2021. "Deontic Logic and Natural Language," in *Handbook of Deontic Logic* Vol. II., Dov Gabbay, John Horty, Xavier Parent, Ron van der Mayden, and Leendert van der Toree (eds), N.p.: College Publications.
[Google Scholar](#) [Google Preview](#) [WorldCat](#) [COPAC](#)
- Cariani, Fabrizio, Magdalena Kaufmann, and Stefan Kaufmann. 2013. "Deliberative Modality Under Epistemic Uncertainty," *Linguistics and Philosophy* 36: 225–259.
[Google Scholar](#) [WorldCat](#)
- Carr, Jennifer. 2015. "Subjective Ought," *Ergo* 2/27: 678–711.
[Google Scholar](#) [WorldCat](#)
- Charlow, Nate. 2018. "Decision-Theoretic Relativity in Deontic Modality," *Linguistics and Philosophy* 41: 251–287.
[Google Scholar](#) [WorldCat](#)
- Chisholm, Roderick. 1963. "Contrary-to-Duty Imperatives and Deontic Logic," *Analysis* 24/2: 33–36.
[Google Scholar](#) [WorldCat](#)
- Chrisman, Matthew. 2012. "'Ought' and Control," *Australasian Journal of Philosophy* 90/3: 433–451.
[Google Scholar](#) [WorldCat](#)
- Dowell, J. L. 2011. "A Flexible Contextualist Account of Epistemic Modals," *Philosophers' Imprint* 11/14: 1–25.
[WorldCat](#)
- Dowell, J. L. 2013. "Flexible Contextualism about Deontic Modals: A Puzzle about Information-Sensitivity," *Inquiry* 56/2–3: 149–178.
[Google Scholar](#) [WorldCat](#)
- von Fintel, Kai. N.d. "The Best We Can (Expect to) Get?," ms. <https://web.mit.edu/fintel/fintel-2012-apa-ought.pdf>
[WorldCat](#)
- von Fintel, Kai and Sabine Iatridou. 2006. "What to Do If You Want to Go to Harlem," ms for the 2005 Rutgers Semantic Workshop.
<https://web.mit.edu/fintel/fintel-iatridou-2005-harlem.pdf>

WorldCat

Forrester, J. W. 1984. "Gentle Murder, or the Adverbial Samaritan," *Journal of Philosophy* 81: 193–197.

[Google Scholar](#) [WorldCat](#)

Gibbard, Allan. 1990. *Wise Choices, Apt Feelings: A Normative Theory of Judgment*. Cambridge, MA: Harvard University Press.

[Google Scholar](#) [Google Preview](#) [WorldCat](#) [COPAC](#)

Jackson, Frank and Robert Pargetter. 1986. "Oughts, Options, and Actualism," *Philosophical Review* 95/2: 233–255.

[Google Scholar](#) [WorldCat](#)

Kaplan, David. 1989. "Afterthoughts," in *Themes from Kaplan*, Joseph Almog, John Perry, and Howard Wettstein (eds), Oxford: Oxford University Press, 565–614.

[Google Scholar](#) [Google Preview](#) [WorldCat](#) [COPAC](#)

p. 502 Kolodny, Nico and John MacFarlane. 2010. "Ifs and Oughts," *Journal of Philosophy* 107: 115–143.

[Google Scholar](#) [WorldCat](#)

Kratzer, Angelika. 1981. "The Notional Category of Modality," in *Words, Worlds, and Contexts: New Approaches in Word Semantics*, Hans Eikmeyer and Hannes Rieser (eds), Berlin: De Gruyter, 38–74.

[Google Scholar](#) [Google Preview](#) [WorldCat](#) [COPAC](#)

Kratzer, Angelika. 1991. "Modality," in *Semantics: An International Handbook of Contemporary Research*, Arnim von Stechow and Dieter Wunderlich (eds), Berlin: de Gruyter, 639–650.

[Google Scholar](#) [Google Preview](#) [WorldCat](#) [COPAC](#)

Kratzer, Angelika. 2012. *Modals and Conditionals*. Oxford: Oxford University Press.

[Google Scholar](#) [Google Preview](#) [WorldCat](#) [COPAC](#)

Lassiter, Daniel. 2017. *Graded Modality*. Oxford: Oxford University Press.

[Google Scholar](#) [Google Preview](#) [WorldCat](#) [COPAC](#)

Lewis, David. 1975. "Adverbs of Quantification," in *Formal Semantics of Natural Language*, Edward L. Keenan (ed.), Cambridge: Cambridge University Press, 3–15.

[Google Scholar](#) [Google Preview](#) [WorldCat](#) [COPAC](#)

Lewis, Karen 2020. "Metasemantics without Semantic Intentions," *Inquiry* 11/14: 1–25. DOI: 10.1080/0020174X.2020.1847184.

[Google Scholar](#) [WorldCat](#)

MacFarlane, John. 2014. *Assessment Sensitivity*. Oxford: Oxford University Press.

[Google Scholar](#) [Google Preview](#) [WorldCat](#) [COPAC](#)

Matthewson, Lisa. 2016. "Modality," in *The Cambridge Handbook of Formal Semantics*, Maria Aloni and Paul Dekker (eds), Cambridge: Cambridge University Press, 525–559.

[Google Scholar](#) [Google Preview](#) [WorldCat](#) [COPAC](#)

Ninan, Dilip. 2005. "Two Puzzles about Deontic Necessity," in *New Work on Modality, MIT Working Papers in Linguistics, #51*, Jon Gajewski, Valentine Hacquard, Bernard Nickel, and Seth Yalcin (eds), 1–27, Cambridge: MIT.

Portner, Paul. 2007. "Imperatives and Modals," *Natural Language Semantics* 15: 351–383.

[Google Scholar](#) [WorldCat](#)

Rett, Jessica 2016. "On a Shared Property of Deontic and Epistemic Modals," in *Deontic Modality*, Nate Charlow and Matthew Chrisman (eds), Oxford: Oxford University Press, 200–229.

[Google Scholar](#) [Google Preview](#) [WorldCat](#) [COPAC](#)

- Rieppel, Michael. 2011. "Stoic Disagreement and Belief Retention," *Pacific Philosophical Quarterly* 92: 243–262.
[Google Scholar](#) [WorldCat](#)
- Ross, Alf. 1944. "Imperatives and Logic," *Philosophy of Science*: 30–46.
[Google Scholar](#) [WorldCat](#)
- Saint-Croix, Catherine and Thomason, Richard. 2019. "Chisholm's Paradox and Conditional Oughts," *Journal of Logic and Computation* 29/3: 369–386.
[Google Scholar](#) [WorldCat](#)
- Schroeder, Mark. 2008. *Being For*. New York: Oxford University Press.
[Google Scholar](#) [Google Preview](#) [WorldCat](#) [COPAC](#)
- Schroeder, Mark. 2011. "Ought, Agents, and Actions," *Philosophical Review* 120/1: 1–41.
[Google Scholar](#) [WorldCat](#)
- Schroeder, Mark. 2015. "Attitudes and Epistemics," in *Expressing Our Attitudes: Explanation and Expression in Ethics*, Vol. II. New York: Oxford University Press, 225–256.
[Google Scholar](#) [Google Preview](#) [WorldCat](#) [COPAC](#)
- Silk, Alex. 2014. "How to Be an Ethical Expressivist," *Philosophy and Phenomenological Research* 91: 47–81.
[Google Scholar](#) [WorldCat](#)
- Silk, Alex. 2016. *Discourse Contextualism*. Oxford: Oxford University Press.
[Google Scholar](#) [Google Preview](#) [WorldCat](#) [COPAC](#)
- Starr, W. 2016. "Dynamic Expressivism about *Deontic Modality*," in *Deontic Modality*, Nate Charlow and Matthew Chrisman (eds), Oxford: Oxford University Press, 355–394.
[Google Scholar](#) [Google Preview](#) [WorldCat](#) [COPAC](#)
- Stojnić, Una. 2021. *Context and Coherence*. Oxford: Oxford University Press.
[Google Scholar](#) [Google Preview](#) [WorldCat](#) [COPAC](#)
- Svavarsdottir, Sigrun. 1999. "Moral Cognitivism and Motivation," *Philosophical Review* 108/2: 161–219.
[Google Scholar](#) [WorldCat](#)
- Wallace, R. Jay. 1998. "Moral Motivation," in *The Routledge Encyclopedia of Philosophy*, Edward Craig (ed.), London: Routledge, 522–528.
[Google Scholar](#) [Google Preview](#) [WorldCat](#) [COPAC](#)
- Wedgwood, Ralph. 2016. "Objective and Subjective 'Ought'," in *Deontic Modality*, N. Charlow and M. Chrisman (eds), Oxford: Oxford University Press, 143–168.
[Google Scholar](#) [Google Preview](#) [WorldCat](#) [COPAC](#)
- Willer, Malte. 2014. "Dynamic Thoughts on Ifs and Oughts," *Philosophers' Imprint* 14/28: 1–30.
[Google Scholar](#) [WorldCat](#)
- Yalcin, Seth. 2012. "Baysean Expressivism," *Proceedings of the Aristotelian Society* CXII/2: 123–160.
[Google Scholar](#) [WorldCat](#)