View metadata, citation and similar papers at core.ac.uk
brought to you by CORE
provided by PhilPapers

DEONTIC LOGIC AND NATURAL LANGUAGE

Fabrizio Cariani

[Forthcoming in *Handbook of Deontic Logic Vol. II*, Dov Gabbay, John Horty, Xavier Parent, Ron van der Meyden, and Leon van der Torre (eds.), College Publications. Comments for minor corrections still welcome!] *

Abstract

There has been a recent surge of work on deontic modality within philosophy of language. This work has put the deontic logic tradition in contact with natural language semantics, resulting in significant increase in sophistication on both ends. This chapter surveys the main motivations, achievements, and prospects of this work.

1 Introduction

The last couple decades have seen a remarkable amount of activity in philosophy of language on the topic of deontic modality. This kind of interaction between philosophy of language and deontic logic is potentially fruitful in both directions. Philosophers of language and linguists leverage their frameworks and techniques to open up new approaches to classic problems in deontic logic. In some cases, problems that appear minor from the perspective of pure logic can become targets of extended analysis in the theory of linguistic meaning. In the opposite direction, the power and sophistication of logical methods can help systematize, constrain and

*I thank Ana Arregui, Bob Beddor, Simon Goldstein, Aynat Rubinstein, Paolo Santorio, and Malte Willer for conversations, feedback and help with tracking down some papers. I also thank an anonymous reviewer for the Handbook of Deontic Logic for extensive comments, as well as the editors for the invitation to contribute as well as feedback on a previous draft.

investigate the space of available answers to linguistic questions. While this second direction of cooperation has received less attention, it seems both possible and desirable for it to become more prominent.[1]

In this paper, I survey some of the most striking contributions in this emerging area. This is the structure of this essay: Section 2 clarifies the linguistic scope of the survey — identifying the expressions and concepts that will be in focus; Section 3 surveys work concerning the interaction between deontic modals and conditionals; Section 4 surveys work on a family of puzzles concerning monotonicity properties of deontic operators in natural language; Section 5 takes up linguistic work on the variety of forces for deontic modals–in particular, work on the difference between *ought* and *must*. Section 6 covers work on the relationship between deontic language and the language of agency. The four sections starting with Section 3 are all independent of each other.

Needless to say, all of these topics have direct antecedents in the deontic logic literature. Since my focus is on interactions with research on natural language, I won't chart all the relevant historical references very carefully. Abundant references to these antecedents are available in the other essays in the present handbook — in particular, Hilpinen and McNamara (2013).

2 Deontic modality as a linguistic category

Let us start with some rough characterizations. What is a deontic modal? We might follow the opening move in Portner (2009) and claim that "modality is the linguistic phenomenon whereby grammar allows one to say things about, or on the basis of, situations which need not be real" (p. 1). This seems right but it is not necessarily a great guide when it comes to classifying specific expressions as modals. I propose for the purposes of this essay to complement it with a sufficient condition for modality: classify an expression e as a *modal* (within a given semantic theory T), if e's semantic evaluation rules manipulate a world of evaluation.[2]

Some theorists operate under a definition that is at the same time more precise, stronger, more theoretically loaded, and ultimately more

[1] Holliday and Icard (2018) make the case for why it's desirable; Holliday and Icard (2017) and Van De Putte (2018) show a possible shape such work might take.

[2] The definition ties the class of modals to one's choice of semantic theory. However, we often talk about modals without specifying a theory. For example, English *may* is generally classified as a modal regardless of one's particular theory. I understand (unrelativized) claims that e is a modal as conveying implicit acceptance that e qualifies as modal according to any reasonable choice of semantic theory.

dubious. According to this definition, all modals express concepts of *possibility* or *necessity*. I do not accept this characterization because I think there are modal expressions that are not well understood as either possibility or necessity operators. To give only one example, probability operators are not well understood as possibility or necessity operators (Yalcin, 2010; Lassiter, 2011, 2017). Additionally, there are analyses of *ought* we will encounter in this essay that deny that *ought* is accurately classified by this scheme. Finally, it is plausible that deontic comparative adjectives are modal in character without being either necessity or possibility operators. Consider for an example the propositional uses of *better* in:

(1) It is better to ride a bike than to drive.

Taken together, these considerations make it preferable to stick with a less committal characterization.

It is an often noted feature of modals that they give rise to a variety of interpretations. Thus, English *may* can convey epistemic possibility (roughly: compatibility with some state of information) or deontic permission (roughly: compatibility with a body of norms). Neither does the list end there: the literature recognizes metaphysical, ability-based, temporal interpretations and many others. That said, not all modals allow every interpretation, and, as we will shortly see, some do not seem polysemous at all.

Deontic interpretations concern modal statuses that obtain in virtue of some norm, or value. Permission and obligation are both prime examples. Following a useful suggestion by Portner (2009), we can see deontic interpretations as part of a broader category of *priority* interpretations of modals.

> "The idea behind the term "priority" is that such things as rules, desires, and goals all serve to identify some possibility as better than, or as having higher priority than, others." (Portner, 2009, p. 135)

This is a felicitous classification suggestion. Looking at the class of priority modals helps zero-in on some generalizations that are harder to detect when we look at broader classes (e.g., modals generally). And not many new generalizations emerge when we look at more specific classes. For example, there are very few general facts about the deontic *ought* that lack matching facts concerning more goal-driven ones.

There is disagreement about what gives rise to this rich variety of interpretations. A view that is often associated with Kratzer's program is that much of the variety of modal interpretations can be traced down to context-sensitivity.[3] While Kratzer's contextualism is superior to a straightforward ambiguity theory, there is a growing intellectual demand for an account of the diversity of interpretations of modals that derives them on the basis of more systematic consideration.

An alternative is to derive the different modal interpretations on the basis of structural facts. Hacquard (2006; 2010) pioneered an approach to do so within a framework that blends elements of Kratzer's approach with event-semantics and emphasizes the syntactic differences between modal sentences carrying epistemic interpretations and ones that carry non-epistemic interpretations.

Many of the questions that are discussed in this survey are, to a degree, independent of one's views on this matter. But no story about deontic modality in natural language is complete without some account of it.

As I anticipated, it is also important to recognize that not all modals are polysemous. Deontic concepts can be expressed by modal verbs like *require*, *obligate*, *permit*, as well as their nominalizations. These verbs do not generally give rise to the broad variety of interpretations that is associated with, say, *must*. In this connection, Hacquard (2013) distinguishes between *grammatical* modals and *lexical* modals. The hallmarks of grammatical modals (like *must*, *ought*, *may*, *can*, etc.) are that they are closed class expressions and are polysemous. (Szabó (2015) suggests that the category of closed class expressions is a useful proxy in natural language for what logicians call "logical constants".)By contrast *lexical* modals (*likely*, *obligatory*, *permitted*) are open class expressions and are typically not polysemous.[4]

The research surveyed here focuses heavily on grammatical modals, and so on the deontic interpretations of *must*, *ought*, and *may*. This is no doubt in part for contingent reasons and in part because the closed class expressions appear closer to the architectural features of a language. Despite that, grammatical priority modals are not insulated. They are

[3]Kratzer (1977, 1981, 1991b, 2012). For rich developments of the contextualist view of modal flavors see also Dowell (2011, 2013); Bronfman and Dowell (2018, 2016).

[4]This distinction has heuristic value but it is unlikely to be perfectly clear cut, since there are evidently lexical modals that do appear to be polysemous. For example, *warranted* and *justified* seem to have both epistemic and deontic interpretation, while *compel* seem to have both causal and deontic interpretations.

embedded in inferential networks that relate them to lexical modals as well. So, our discussion will occasionally touch on lexical modality.

3 From conditional obligations to iffy oughts

A historically important tradition in deontic logic focuses on the concept of *conditional obligation*—what one is obliged to do given that some condition holds. Philosophers of language and linguists have revised and remixed some of the main arguments in this tradition. Their distinctive concerns have led to substantial progress, as well as to the opening of some new avenues of inquiry.

Let us take off our exposition from a famous passage in David Lewis's "Semantic Analyses for Dyadic Deontic Logic" (Lewis, 1974). Keep in mind that, at this point in time, Lewis is already summarizing a wealth of prior work on conditional obligation.

> "It ought not to be that you are robbed. A fortiori, it ought not to be that you are robbed and then helped. But you ought to be helped, given that you have been robbed. The robbing excludes the best possibilities that might otherwise have been actualized, and the helping is needed in order to actualize the best of those that remain. Among the possible worlds marred by the robbing, the best of a bad lot are some of those where the robbing is followed by helping." (Lewis, 1974, p.1)

Lewis took these considerations to motivate an analysis of dyadic conditional obligation. As will no doubt be familiar, standard deontic logic is built on the idea that obligation operators are unary necessity operators. Following an extended analogy with conditional probability, we might step away from that paradigm and instead introduce *binary* obligation operators. For instance, the operator $\bigcirc(\mathsf{Q} \mid \mathsf{P})$ could be used so as to mean that Q is obligatory given that we restrict focus to the possibilities satisfying P. Lewis's example could be modeled as follows: $\bigcirc(\textit{you are helped} \mid \textit{you are robbed})$.

Studying the logic of such dyadic operators is an entirely unobjectionable enterprise, especially if it is divorced from considerations of natural language. However, many authors have pointed out that, from the point of view of our understanding of natural language, it is a mistake to think of *iffy oughts* like (2) as having the form $\bigcirc(\textit{nap} \mid \textit{tired})$.[5]

[5] I borrow the 'iffy ought' terminology from Willer (2012).

(2) If Iris is tired, she should take a nap.

One reason for resisting the representation in terms of conditional obligation operators is that it misses out on the generality of conditional modality. For one thing, we may want to analyze sentences with multiple operators in their consequent, such as:

(3) If Iris is tired, she will try to stay awake but she should take a nap.

For another, each deontic modal expression would need its own dyadic operator. After all, we can have conditional permissions as in (4) and other kinds of conditional deontic claims as in (5).

(4) If you have the permit, you may fix your sidewalk.

(5) If you bought this guitar, you must buy this amplifier.

Setting generality aside, from the point of view of theory-design, the idea of dyadic obligation operators seems insufficiently modular. If the point of the formal semantic system is to contribute to a compositional theory of meaning for a language like English, it should separate the contributions of *if* and *ought* (*viz.*, *if+may* and *if+must*, depending on the case). As Thomason put the point:

> "A proper theory of conditional obligation [...] will be the product of two separate components: a theory of the conditional and a theory of obligation." (Thomason, 1981b, p.165)

Here is Bonevac riffing on Thomason's theme:

> "At the very least, a theorist using a conditional obligation operator owes us an explanation of how the semantics of the operator depends on the semantics for obligation and the conditional simpliciter. Sentences expressing conditional obligations are intelligible to anyone understanding *should* (or *ought to*) and *if*. The combination of these words is no idiom. The meanings of such sentences, therefore, should be explicable in terms of the meanings of *if* and *should* construed independently." (Bonevac, 1998, p.37)

Thomason and Bonevac's comments seem entirely right to me. Anyone who cares about a formal theory of meaning should heed their advice and reject conditional obligation theories. The next question then is what *do* we learn by separating out the contributions of *if* and *should*?

3.1 Kratzer's theory of conditionals and modals

Proponents of conditional obligation operators lost a battle on the interpretation of deontic conditionals like (2), but they ended up scoring an unexpected win in a war they might not even have known they were fighting.

The semantic framework for modality that has come to be viewed as paradigmatic, Kratzer's,[6] incorporates some of the key ideas of the dyadic analysis—while also heeding Thomason's demand for a compositional analysis of sentences like (2). In light of its prominence, and of the fact that much of the literature reacts to this paradigm, I will present a simplified version of Kratzer's theory.[7] I will refer to this as the *baseline* theory.

According to the baseline theory, every modal (and crucially, the grammatical modals *ought*, *must*, *may*) takes two propositional arguments, a restrictor and a scope.[8] The scope proposition is the proposition whose modal status we are interested in. The restrictor proposition delimits the set of worlds that are to count as relevant. Thus in (6), the restrictor is the proposition that you have a permit, the scope is the proposition that you fix your sidewalk.

(6) *may*[you have a permit][you fix your sidewalk]

The aggregate interpretation is something like what we'd express by saying *if you have a permit, you may fix your sidewalk*. Crucially, even the overtly unconditional *you may fix your sidewalk* gets assigned a logical form like:

(7) *may*[RESTRICTION][you fix your sidewalk]

In this representation, RESTRICTION is a place-holder for a restrictor proposition. Restrictor propositions are provided by context, and usually are a bit more generic than the kind of antecedent in (6). Plausibly (7) is interpreted relative to a restriction to the worlds in which the city laws and the relevant circumstances are about the same as they are in the base world. The crucial point is that, according to the baseline analysis, restrictor propositions are *always* needed to interpret modal claims. So far so good: the dyadic obligation theorist has been nodding all along.

[6] Kratzer (1977, 1981, 1991a,b, 2012).

[7] It should be noted that Kratzer has explored, advocated and developed a variety of other frameworks and revisions of her theory.

[8] Modals might, in fact, take more than two arguments depending on the details of the semantic theory. I ignore this complication.

Kratzer adds a few important theses. To start, the theory is to apply to *all* modals, and not limited to some specific interpretation (Kratzer uses "flavors" to refer to what I have described as "interpretations" of modals—i.e., epistemic, deontic, etc.). Furthermore, as I noted in the previous section, these different interpretations arise as a result of different settings for various contextual parameters.

These parameters include the parameters that fix the contextual restrictor proposition. Critically, however, the restrictor proposition may also be affected by explicit elements of the sentence, and perhaps even elements of the discourse at large. The principal (but not the *only*) device of restriction is the conditional *if*. Let us see how all of these ideas play out in the formalism.

I said that the restrictor proposition is partly determined by context. In the baseline theory, context provides values to two parameters:

- a *modal base* f (mapping the set of all worlds Ω to $\mathcal{P}(\Omega)$)
- an *ordering source* g (mapping Ω to preorders $\Omega \times \Omega$)

If fed a world of evaluation w these functions output a set of worlds M (the *modal background relative to* w) and a pre-order of worlds $\prec$.[9]

The concepts of modal base and ordering source help provide lexical entries for modals. These entries are easiest to state under restrictions that entail the *limit assumption* (see below for the content of this assumption). For example, assume that there are finitely many possible worlds. Then we can give the following analyses for the modals *must*, *might* and *ought*. (We momentarily assume that *ought* and *must* have the same meaning, but we will question this assumption in Section 5.)

$$\textsc{best}(f,g,w) = \{u \in f(w) \mid \neg\exists v, v \succ_{g,w} u\}$$

(8) $$[\![ought]\!]^w = [\![must]\!]^w = \lambda\mathbf{P}.\lambda f.\lambda g.\forall v \in \textsc{best}(f,g,w), v \in \mathbf{P}$$

(9) $$[\![might]\!]^w = \lambda\mathbf{P}.\lambda f.\lambda g.\exists v \in \textsc{best}(f,g,w) \ \& \ v \in \mathbf{P}$$

[9]Kratzer's official theory is formulated in the framework of premise semantics. This means that the types of these contextual parameters are slightly different from what I have suggested. In particular, modal bases are actually functions from worlds to sets of propositions ($\Omega \mapsto \mathcal{P}(\mathcal{P}(\Omega))$) which then determine a set by intersection. Similarly, her ordering sources are also functions from worlds to sets of propositions, which then determine a pre-order. The premise semantics formulation of the theory has some explanatory advantages that will come in handy in Section 5. But for now it will be quicker to set it aside.

To illustrate these clauses, consider evaluating *you must run* at world w. This will be true (at w) if every world in $\textsc{best}(f, g, w)$ is a world in which you run. The worlds in $\textsc{best}(f, g, w)$ are exactly those worlds in $f(w)$ such that no world is better than them (in the sense of $\prec_{g,w}$). If $f(w)$ contains a world in which you walk that is better than every world in which you run, our target sentence will be false. It will also be false if for every run-world there is a better walk-world (whether or not there is a walk-world that beats all the run-worlds). Otherwise, it'll be true.

The entries above presuppose the infamous *limit assumption*. Part of the content of the limit assumption is the claim that, for every choice of $f(w)$, there are maximal points in the partial order $\geq$.[10] However, it seems plausible, given the deontic interpretation of $\geq$, that this property could sometimes fail to hold of a deontic ordering. For example, there could be setups that contain infinite chains of worlds with monotonically increasing value. To address that possibility, Kratzer proposes a somewhat more complicated quantificational condition.

$$\textsc{forceable}(\mathbf{P}, f, g, w) = \{u \in f(w) \mid \exists v \in f(w), v \geq_{g,w} u \ \& \ \forall z \in f(w), \textit{if } z \geq_{g,w} v, z \in \mathbf{P}\}$$

(10) $$[\![ought]\!]^w = [\![must]\!]^w = \lambda \mathbf{P}.\lambda f.\lambda g.\forall v \in f(w), \textsc{forceable}(\mathbf{P}, f, g, v)$$

In words, $\mathbf{P}$ is forceable from u's perspective if there is a relevant world v that is at least as good as u such that all of the relevant worlds that are at least as good as v make $\mathbf{P}$ true. And *must*(P)/*ought*(P) is true if the prejacent proposition $\mathbf{P}$ is forceable from the perspective of any relevant world. [11]

The limit assumption is an excellent example of a useful idealization. It is almost certainly false, given the intended interpretation of the ordering. But, in nearly every application, working under more realistic

[10] I take it that the content of the limit assumption is not exhausted by this condition. See Kaufmann (2017) for an extensive study, detailing and resolving much outstanding confusion about how the limit assumption ought to be formulated.

[11] Let us consider why (10) and (8) are equivalent in the special case in which there are finitely many worlds. (Recall that this is not the only hypothesis compatible with the limit assumption, but it will be illustrative regardless of that.) Suppose *must* P meets the truth-condition in (8); then all the worlds such that nothing is better than them are P-worlds. But so, if we consider an arbitrary relevant world v we should be able to find whichever world that is identical to or better than v that is maximal and let that world be the witness to the existential quantifier in the definition of FORCEABLE. In the opposite direction, suppose that the proposition expressed by P is FORCEABLE relative to f and g and any relevant world. Consider an arbitrary world v belonging to $\textsc{best}(f, g, w)$, then v is in $f(w)$ so P is forceable with respect to v. However, since v is a terminal world with respect to $\succ_{g,w}$, all that forceability entails is that P must be true at v as demanded by the analysis.

assumptions increases the cognitive load for modest benefit. In general, it is practical to default to working under conditions that support the limit assumption.

The baseline theory of modals is complemented, and completed, by Kratzer's theory of conditionals. As noted, in Kratzer's semantics, *if* does not denote a binary connective. It serves instead to further restrict the modal base of a modal in its scope. It is not straightforward to give a fully compositional implementation of this idea,[12], so we will keep it at a relatively intuitive level. The *effect* of this idea is that sentences that look like (*if* P)(◯Q) are evaluated by:

(i) interpreting the restrictor proposition (i.e. whatever proposition **P** is expressed by P in context);

(ii) restricting the modal base f for ◯ with **P** (i.e. creating a new function $f + \mathbf{P} = \lambda w.f(w) \cap \mathbf{P}$);

(iii) evaluating the modal claim, ◯**Q** relative to this shifted modal base $f + \mathbf{P}$ and whatever ordering source g was provided by the initial context.

Of course, some *if*-sentences do not have overt modals. In these cases, Kratzer's hypothesis is that *if* restricts a *covert* modal. This covert modal defaults to an epistemic necessity interpretation. In other words (11) is actually interpreted as (12) with the restrictor proposition narrowing down the modal base of □.[13]

(11) If she called, she lost.

(12) If she called, □ (she lost).

With the baseline view developed let us see how it applies to some classic problems from deontic logic.

3.2 Chisholm's paradox

The baseline theory makes a distinctive prediction about the classical paradoxes of conditional obligation. I illustrate the prediction by considering Chisholm's paradox (Chisholm, 1963). The theory makes similar

[12]See von Fintel (1994, ch.3) for an example of the sort of work that is involved.

[13]This opens up the possibility that there might also be cases in which a covert modal is posited in addition to an overt one. This possibility is leveraged in Geurts (unpublished) and discussed further with application to conditional deontics in Cariani *et al.* (2013).

predictions in the case of the Good Samaritan paradox and Forrester's Gentle Murder paradox (Forrester, 1984). I will then compare the baseline with some alternatives.

Here is a slightly touched-up version of Chisholm's familiar vignette:

> Your elderly neighbors asked for your assistance to prepare their taxes. Because you are their only acquaintance, *you ought to help them*. But, they are easily frightened by potential intruders. So *it ought to be that, if you go, you tell them in advance*. However, telling them that you are going will be bad if you are *not* going. So, *if you don't go, you ought not to tell them*. As it happens, your favorite show is on TV and *you don't go*.

The four italicized sentences generate the puzzle. To get a little closer, we can do a preliminary formalization of these:

(i) $\bigcirc(go)$

(ii) $\bigcirc(\textit{if } go, \textit{ tell})$

(iii) $\textit{if } \neg\, go, \bigcirc(\neg\, \textit{tell})$

(iv) $\neg go$

This formalization is not meant to be definitive. Depending on our background commitments, we might want to revise aspects it. For example, defenders of the baseline theory might want to revise (ii) in ways that will become clear shortly.

Whatever the final formalization, critical elements in the paradox are already apparent at this preliminary level. Chisholm's key assumption is that $\bigcirc$ scopes outside the conditional in (ii) but inside the conditional in (iii).

Chisholm noted that these sentences sound collectively consistent. The puzzle is that they fail to be consistent under standard assumptions of deontic logic, combined with two principles governing the interaction of conditional and obligation operators. The two principles are "factual" and "deontic" detachment. It is indeed common (though perhaps not entirely correct) to frame Chisholm's puzzle as pitching these against each other (Loewer and Belzer, 1983).

FACTUAL DETACHMENT (FD). $\mathsf{P}, (\textit{if } \mathsf{P},\ \bigcirc\mathsf{Q}) \vdash \bigcirc(\mathsf{Q})$

(WIDE) DEONTIC DETACHMENT (DDW). $\bigcirc(\mathsf{P}), \bigcirc(\textit{if } \mathsf{P}, \mathsf{Q}) \vdash \bigcirc(\mathsf{Q})$

Though I formulate deontic detachment with $\bigcirc$ taking with scope over a conditional, it is also worth considering a version of deontic detachment in which the deontic modal appears in the consequent of the conditional.

NARROW DEONTIC DETACHMENT (DDN). $\bigcirc(\mathsf{P}), (\textit{if}\ \mathsf{P}, \bigcirc\mathsf{Q}) \vdash \bigcirc(\mathsf{Q})$

After all, a defender of the restrictor analysis might formulate the second premise of Chisholm paradox as: *if go*, $\bigcirc(tell)$

Chisholm's formulation of the paradox indicates that he intends to construe premise (ii) with the conditional scoping under the deontic modal. That would give relevance to DDW as far as the argument goes. However, as Saint Croix and Thomason (2014) forcefully note, there are very few constructions in English that are appropriately represented with a deontic modal scoping over a conditional. In the following, I use DD for points that go through on either construal of this principle.

Under the assumption that sentences of the form (*if* P, $\bigcirc$Q) are conditionals with $\bigcirc$ in the consequent, FD just is *modus ponens*. As we saw, this assumption is strictly speaking false on the baseline view, since on the baseline view conditionals restrict modals, as opposed to *connecting* pairs of propositions. To minimize our committments, it is be preferable to stick to calling it "factual detachment".

Here is a sketch of the proof of inconsistency, taking the four sentences as premises:

(v)	$\bigcirc(tell)$	(i),(ii) DD
(vi)	$\bigcirc(\neg tell)$	(iii),(iv) FD
(vii)	$\bigcirc(tell\ \&\ \neg tell)$	(v),(vi) AGGLOMERATION
(viii)	$\bot$	(vii), AXIOM D

There are many familiar options to address the paradox. Giving up deontic detachment, giving up factual detachment for *iffy* oughts; rejecting agglomeration;[14] giving up some structural rules; arguing that the premises need to be formalized differently. The question that matters for present purposes is which of these answers best fits with the constraints and commitments coming from various theories within natural language semantics.

[14] If we assume the duality of obligation and permission, the agglomeration step could be avoided in the proof. We might reason from (vi) to $P(\neg tell)$ via axiom D which would directly contradict $\neg P \neg tell$, which is equivalent to (v). In general, the strategies I will mention below that block the agglomeration step, must also reject duality.

For example, I mentioned above that the baseline theory makes a distinctive prediction. In particular, the theory allows all four sentences to be true at once by blocking FACTUAL DETACHMENT. To see this, model the case by assuming that there are four worlds in the modal base $w_{GT}, w_{G\overline{T}}, w_{\overline{G}T}, w_{\overline{GT}}$.[15] Suppose that, as is plausible, these are ranked as follows in the relevant ordering source:

$$w_{GT} > w_{G\overline{T}} > w_{\overline{GT}} > w_{\overline{G}T}$$

This parameter setting has the following effects. The unconditional claims $\bigcirc(go)$ and $\bigcirc(tell)$ are true because w_{GT} is the best world. For the same reason, $\bigcirc(\neg tell)$ is false. If the second premise is interpreted along the lines of DDN, it is easy to see that the conditional antecedent restricts the domain to the two top-most worlds. Among these, the best is the one in which you tell.[16] For the third premise, processing the conditional antecedent *if you don't go* restricts the modal background to $\{w_{\overline{GT}}, w_{\overline{G}T}\}$ world. The set of best worlds in this restricted background is the singleton $\{w_{\overline{GT}}\}$. So, under the restriction to $\neg go$, $\bigcirc(\neg tell)$ is true. For the fourth premise, just assume that the actual world is either of the $\overline{G}$-worlds.

The upshot: the baseline validates DD, AGGLOMERATION, AXIOM D but invalidates FD.

One worry about this type of approach, articulated in Arregui (2010), is that part what's puzzling about Chisholm's case can be replicated without at all involving FD. Indeed, we have a lingering, and hiterto unaddressed, intuition that the relevant obligations one is under might change with time, and specifically might depend on how some things turn out in the actual world. Once it is settled that you won't go to help your

[15] The indices are designed to convey which of *go* and *tell* are true at each world.

[16] Things are a bit more complicated if we formalize the premise as DDW. If we have to stick to the idea of wide-scoping $\bigcirc$, the baseline predicts the following logical form for premise (ii):

$$\bigcirc(\textsc{restriction})(if\ go, \Box tell))$$

In this, RESTRICTION picks out the contextual restriction (which in this case we assume to be tautological, so that no worlds are ruled out). Whether premise (ii) is true, will depend on what kind of modal $\Box$ is. If its domain at each world is simply $\{w_{GT}, w_{G\overline{T}}, w_{\overline{G}T}, w_{\overline{GT}}\}$, then the sentence will be false. It could be made true by forcing the modal $\Box$ to have stricter domains. For example, the modal base and ordering source of $\Box$ could be set up so that $if\ go, \Box tell$ is equivalent to the material conditional $go \supset tell$. This could be accomplished for instance by requiring that for each world w, w is the unique highest ranked modal in the ordering source of $\Box$ (even if it isn't the highest ranked world in the ordering source of $\bigcirc$). Simulating the material conditional within the baseline view approximates the interpretation that Chisholm had in mind for sentence (ii) and also makes it true relative to our chosen parameter settings.

neighbor, it is perhaps no longer the case that you ought to go. The point is best illustrated with some nearby cases. Following Arregui, consider the argument from (13)-(14) to (15):

(13) It ought to be the case that Nina does her job well.

(14) It ought to be the case that if Nina does her job well, she gets a promotion.

(15) It ought to be that Nina gets a promotion.

Surely whether Nina gets a promotion does not just depend on the purely normative premises (13) and (14). Whether she gets a promotion should also depend on how things are in the actual world—for example, on whether she does her job well.[17]

The theoretical proposal of Arregui (2010) involves a solution to Chisholm's paradox that involves giving up *both* DD and AGGLOMERATION. I will present Arregui's semantics in some detail, for it is extraordinarily interesting and somewhat neglected in the philosophical literature.

Arregui presents her preferred semantics for $\bigcirc$ within the framework of situation semantics.[18] Start with the idea that possible worlds have parts. Call these parts *situations*. Assume that situations are structured by a relative parthood relation which is reflexive, antisymmetric and transitive. Possible worlds are maximal situations with respect to the parthood relation. Say that a situation s *deontically requires* P if all the deontically best ways of extending s into a possible world make P true. My promising to help you requires my helping you: this is because in every ideal world that extends the situation of the promise I help you.

Now, suppose that you want to evaluate $\bigcirc$(P) at world w. Consider any situation s that is (i) part of w (ii) compatible with P and (iii) not enough to guarantee P on its own. Say that these are the *P-relevant* situations in w.[19] With these concepts in hand we can state Arregui's semantic theory:

(16) $\bigcirc$(P) is true at w if every s that is P-relevant in w is part of a situation s' that requires P.

[17] See also (Lassiter, 2017, §8.6) for a discussion that echoes these points.

[18] I don't think this is *essential* but it certainly is heuristically helpful.

[19] Arregui's situations are world-bound, in the sense that they belong to at most one world, though they may have counterparts in other worlds.

Note the double layer of quantification: every P-relevant situation has to be extendable to a situation s' that requires P.[20]

We can illustrate these ideas by considering their application to (15) — *It ought to be that Nina gets a promotion*. This is predicted to be false in a world in which Nina does not do her job well and true in a world in which she does. To see why it's false in worlds in which Nina does not do her job well, let s be a situation that includes her accepting the job and every instance of Nina doing the job poorly. Intuitively, there is no way of extending this situation into one that requires that she be promoted. To see why it's true in worlds in which Nina does does her job well, start with any situation s, and extend s to a situation s' that includes the situation of Nina doing her job well. It is intuitive that such a situation s' deontically requires her getting the promotion—and plausibly this requirement will be formalized once we have a correct theory of premise (14).

For conditional oughts, Arregui proposes an account that works out to the following truth-conditions. To evaluate (*if* P)($\bigcirc$ Q) at w, consider all the P-relevant situations s in w; the claim will be true if every such s can be extended to an s' that in turn can be extended to a minimal P-satisfying s'' such that s'' requires P. For (14), imagine extending Nina's accepting the offer to a situation s'' in which she does her job well, then s'' requires Nina's getting a promotion, making the conditional true.

It is an important feature of this theory (though not one that is explicitly discussed in Arregui's paper) that it invalidates both DD and AGGLOMERATION of $\bigcirc$ over conjunction. That is, the proof in (v)-($viii$) contains two invalid steps (given the semantics), namely (vi) and (vii). We might find it surprising that the solution is not minimal. If Chisholm's paradox is to drive us away from some ideal of classicality, some might prefer to be driven as little as possible.

However, it turns out that the non-minimality of Arregui's solution to Chisholm's paradox is a crucial feature, and not a bug. This is illustrated by another recent observation in the literature: no completely satisfactory solution to Chisholm's paradox can be limited to abandoning DEONTIC DETACHMENT. This was initially pointed out by Saint Croix and Thomason (2014), and later echoed by Fine (unpublished). These authors show that

[20] A moment's reflection should clarify why the concept of P-relevance is needed: situations incompatible with P are not extendable into possible worlds that verify P, whether ideal or not; and the situations that guarantee P are trivially extendable.That said, omitting condition (iii) would probably not affect the truth conditions. Since (16) quantifies universally over P-relevant situations, it'll do no harm to include some that are trivially part of situations that require P.

there are versions of Chisholm's paradox that do not involve DD at all. Consider this:

> You ought to go help your elderly neighbors. But if you don't help them, you ought to apologize to them. However, your favorite show is on TV and so you don't go.

From the second and third premise, derive that you ought to apologize via FD. From this, via AGGLOMERATION, derive that you ought to go *and* apologize. While that is not itself contradictory, it is a rather odd consequence that involves only FD and AGGLOMERATION. In sum, giving up DD cannot be the entire story.

The accounts of Chisholm's paradox we have reviewed until now require giving up at least one of FD or DD. But another angle that has received substantial exploration in the philosophy of language literature concerns the prospects for an account of Chisholm's paradox that gives up *neither* of these. Views with this feature have been explored independently and with very different tools by Willer (2014) and Saint Croix and Thomason (2014), in philosophy of language. These works parallel a related tradition that explores the same conceptual space within deontic logic.[21]

The essence of both these approaches is that the fourth premise (*you don't go*) changes the context in a way that undermines the first (*you ought to go*). These views accept that $\bigcirc(\neg tell)$ does (in some sense) follow from (iii) and (iv) and $\bigcirc(tell)$ does (in some sense) follow (v) and (vi). Modeling the context-shift allows us to account for why all four sentences seem acceptable while denying that they all are acceptable in one and the same context. The two models differ in how they represent the context shift and in what relations they bring to bear to the account of inference patterns.

Willer models this insight in the context of update semantics (extending work by Veltman (1996) to the deontic case): a very rough way of putting his view is that he thinks that (a) FD and DD are both valid; (b) however valid patterns of inference are defeasible. In the specific case, the application of DD at step (v) is disallowed because the conditional *ought* in premise (ii) is defeated by the information in premise (iv), i.e. that you don't go.

By contrast, Saint-Croix and Thomason develop their account in a more standard contextualist static semantics. Strictly speaking, in fact,

[21] Prakken and Sergot (1996, 1997); Jones and Pörn (1985); Carmo and Jones (2002); Gabbay (2012).

Saint-Croix and Thomason do not *validate* DD. Instead, on their proposal it turn out to be "pragmatically valid" roughly in the sense in which Stalnaker (1976) talks about "reasonable inference". One may reasonably expect that natural language inferences are not accounted for purely in terms of an unaided concept of logical validity. Instead, we supplement our arsenal with a pragmatic notion that tracks something like preservation of acceptability. What Saint-Croix and Thomason suggest is that DD (and in particular DDN) meets this criterion.

To wrap up this discussion, I want to draw attention to some general morals about the significance of this debate, rather than probe the individual views. The important point, which seems undeniable to me, is that linguistically oriented investigation into Chisholm's paradox has substantially deepened our understanding of the paradox. We have a better sense of the constraints on a solution to it and of the tradeoffs involve in various approaches. In particular, the collective moral of the papers I have just surveyed is that a complete solution to Chisholm's paradox must at least:

(1) explain why some *ought*-claims seem to depend on what is actual,

(2) explain why contrary-to-duty oughts sometimes cannot be conjoined felicitously with unconditional ones, and

(3) be faithful to the particular way in which *ought*-sentences are context-dependent.

In addition, this investigation has improved our understanding of Kratzer's semantics. It is only as part of work like this that researchers have started to emphasize the fact that Kratzer semantics invalidates *modus ponens* — or at least, what on the surface looks like *modus ponens*. Applying theories of the linguistic meaning of modals to well-known problems in logic turns out to be a remarkable way of gaining insight about both.

3.3 The miners paradox

Philosophers of language have also infused fresh blood into the discussion of the *miners paradox*. While the history of the miners paradox places it more squarely within the tradition of metaethics rather than that of deontic logic,[22] the paradox does have implications for deontic logic, and particularly so in the revamped version that has become the center of so much recent research.

[22] Previous work that is of significance here is Regan (1980); Parfit (unpublished); Jackson (1991); Goble (1996).

This resurgence started with an influential article by Kolodny and MacFarlane (2010). They represent the miners scenario as follows:

> "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, the lowest in the shaft, will be killed."

The paradox is this: intuitively, all of the following sentences seem acceptable in the given scenario.

(17) Either the miners are in shaft *A* or in shaft *B*.

(18) If the miners are in shaft *A*, we should block shaft *A*.

(19) If the miners are in shaft *B*, we should block shaft *B*.

(20) We should block neither shaft.

The joint acceptability of these sentences suggests that they ought to be consistent, given a semantic theory for *ifs* and *oughts*.[23] However, they turn out to be inconsistent under basic assumptions. In particular, they are inconsistent if *if* and *or* are both interpreted along the lines of classical logic. And, perhaps more importantly, they are inconsistent if the sentences are modeled in the context of the Kratzerian baseline view (even if its account of *if* is non-classical in the ways I have outlined above).

Kolodny and MacFarlane put abundant effort into rejecting the tempting, but mistaken, idea that the Miners Paradox is easily resolved by distinguishing between subjective and objective oughts. They are right that it is not.

More positively, they offer a two-part diagnosis for what goes on in the Miners Paradox. The first part of the diagnosis is that the paradox forces us to invalidate *modus ponens*. The second is that the paradox requires us to adopt a *seriously information sensitive* analysis of *ought* and *should*. (I will explain shortly what this means.)

[23]The data-point is sometimes presented as the claim that all of (17)-(20) are *true*. Since there are views on which *ought*-sentences and conditional sentences do not strictly speaking have truth-conditions (in the sense that they do not divide the worlds in which they are true from the worlds in which they are false), I prefer the less commital terminology of consistency and joint-acceptability. These are concepts that even a non-truth-conditional theorist would need to have a theory of (though perhaps a non-classical theory).

The *modus ponens* diagnosis is widely viewed as off the mark for two reasons. First, it suggests that according to received views *modus ponens* is valid. This is not so on the Kratzerian baseline view (which has as good a claim as any other view to the title of "received view"). As we saw in the discussion of Chisholm's paradox, the baseline invalidates *modus ponens*. However it still makes the foursome inconsistent (as was independently pointed out in Charlow (2013) and Cariani *et al.* (2013)).

More importantly, invalidating *modus ponens* isn't essential to the solution of the Miners Paradox. There are sensible theoretical packages that hold on to *modus ponens* and instead give up some of the inference patterns that depend on subarguments. For example, one can give up the combination of *disjunctive syllogism*, *modus tollens* and *reductio ad absurdum*.[24] The upshot of these two considerations is that the validity of *modus ponens* doesn't really matter all that much to an understanding of the Miners Paradox.

What does matter is the point about information sensitivity. Recall how I described the baseline algorithm for evaluating claims of the form *if* $\mathsf{P}, \bigcirc\mathsf{Q}$. Very roughly:

(i) Ask context for the modal background $f(w)$ of $\bigcirc$ and for some kind of preorder $\succeq_w$ of worlds by their relative priority.

(ii) Use these elements to compose a pre-domain d consisting of the best (according to $\succeq$) worlds in M.

(iii) Evaluate P and compose the final domain D by restricting d with $\mathbf{P}$ (i.e. $D = d \cap \mathbf{P}$).

(iv) Check that all worlds in D verify Q.

The Miners scenario illustrates a defect in this algorithm. The problem is that the ordering of worlds should be sensitive to the shifts that are introduced by conditional supposition. It may be unconditionally best to block neither shaft. But if the miners are in shaft A, it will be best to block A; and if they are in shaft B it will be best to block B. This is not allowed by the baseline algorithm: if w is the best world in some initial information state, then it remains best under any supposition P that is true at w (Charlow, 2013; Cariani *et al.*, 2013).

[24] Various facets of this point are highlighted in Willer (2012), Yalcin (2012), Bledin (2015) Bledin (forthcoming). If the trio of principles above seems to be a jerry-rigged collection, note that they are exactly those principles that are typically associated with local assumptions and subproofs in natural deduction systems.

What makes the semantics information sensitive is that it satisfies a strengthening of the general principle that what one ought to do in an initial information state i, need not be what one ought to do under information states stronger than i. The condition is most simply put in terms of domains of quantification (though the idea of an information sensitive semantics is not limited to quantificational accounts of *ought*). Any quantificational semantics for *ought* should provide some kind of function, call it $\mathtt{best}$ that given an information state i and whatever other general parameters π outputs a domain for *ought*-sentences. The core claim of a quantificational information sensitive semantics is this:

INFORMATION SENSITIVITY it is not the case that if i^+ refines i and i^+ is consistent with $\mathtt{best}(i,\pi)$, then $\mathtt{best}(i^+,\pi) \subseteq \mathtt{best}(i,\pi)$

The bug in the baseline domain determination algorithm is that it only allows the set of ideal worlds to change when the new information rules out *all* of the formerly best worlds.

As for satisfying INFORMATION SENSITIVITY, there are now a plethora of proposals, all of which are distinguished by their own set of advantages and disadvantages.[25] One question in this domain that should perhaps receive more attention is how to reconcile the idea that subjective oughts are sensitive to information with the rather common idea within semantics that deontic modals take *circumstantial* modal bases—that is, modal bases whose restriction is given by facts, or circumstances, in the base world that are possibly beyond the epistemic reach of any of the agents in context.[26]

The debate on information sensitivity has branched out in important ways, beyond the initial discussion of the Miners paradox. Here are a few highlights.

(i) Willer (2016) draws important connections between the information dependence of deontic modals and the non-monotonicity that is represented in frameworks such as update semantics.

(ii) While Kolodny and MacFarlane stopped at qualitative suppositions, Cariani (2016b) argues that we ought to extend the idea of information sensitivity to probabilistic suppositions—e.g. antecedents like

[25] In addition to Kolodny and MacFarlane (2010), see also Cariani *et al.* (2013); Cariani (2016b); Carr (2015); Charlow (2013, 2016); MacFarlane (2014); Silk (2014); Willer (2012, 2016), as well as the related work by Bledin (forthcoming).

[26] The inadequacy of circumstantial modal bases to capture the relevant readings is discussed at some length in Cariani *et al.* (2013), but I have not seen work that has attempted to retrieve the explanatory goodies that the assumptions of circumstantial modal bases does in Kratzer semantics.

if the miners are likely to be in shaft A. Cariani (2016b) develops a system that accommodates probabilistic information sensitivity.

(iii) Defenders of the baseline semantics have identified complex contextualist maneuvers to account for the joint acceptability of (17)-(20). The best iterations of this program von Fintel (2012), Dowell (2013), Bronfman and Dowell (2016) look like they can play to an empirical draw with the best information sensitive approaches. These solutions trade on an important element of the semantics that is often neglected in information sensitive systems: namely, one could represent facts about an agent's information state as included in the world that is given as input to the modal base and ordering source function.

On the other hand, one's mileage may vary on how important it is to hold on to the baseline semantics, especially considering that many of the alternatives are direct generalizations of it. It certainly seems to some theorists (including myself) that the lack of serious information dependence is not just a local empirical defect to be patched with auxiliary assumptions, but a conceptual flaw.

(iv) The idea of an information sensitive semantics can be interestingly extended to other kinds of modals. Lassiter (2011, 2017) and Finlay (2014) independently discuss the case of evaluative adjectives such as *good* and its comparative *better*. Lassiter, developing a suggestion from Levinson (2003), also notes the nearby case of desire attributions. For further inroads in this direction, see the hot-off-the-press Jerzak (forthcoming).

The linguistic territory surrounding the Miners Paradox is well charted at this point.[27] However, there may be some unexplored questions in the context of a theory of normative uncertainty.

It seems plausible to expect that a theory of normative uncertainty would include a model of how we might ascribe degrees of belief to *ought*-claims, and normative claims generally. An initial thought could

[27]In addition to the questions about information sensitivity, the Miners Paradox has also spawned an important discussion on the philosophy of semantics. The question at the center of this thread of research is to what extent we are allowed to incorporate substantive ethical assumptions within deontic semantics. Neutralists have suggested that deontic semantics ought to be as free as possible of substantive ethical assumptions. They include Cariani (2014, 2016b); Carr (2015); Charlow (2016, 2018). For an approach that more boldly incorporates substantive assumptions, see Lassiter (2016) and (Lassiter, 2017, ch.8).

be that these degrees of belief must be *probabilities*. This would extend the standard probabilist tenet that degrees of belief for ordinary factual claims ought to respect the axioms of the probability calculus. Indeed, those who don't see a principled distinction between the information conveyed by factual claims and the information conveyed by normative claims almost certainly have to accept this extension.

However accepting the extension puts us in dangerous territory because of its proximity to Lewis-style triviality (see Charlow (unpublished) for an articulation of how triviality results might extend to deontic claims). Without getting too deep into a presentation of triviality, we can use the Miners paradox to highlight that there are parallel intuitions about credences. If these intuitions are taken at face value, they collectively give rise to probabilistically incoherent credences. Consider these questions, again on the background of the miners' scenario:

- what credence should one assign to the claim that we ought to block shaft A given that the miners are in A?
- what credence should one assign to the claim that we ought to block shaft B given that the miners are in B?
- what credence should one assign to the claim that we ought to block neither shaft?

Under the assumptions in the scenario, it seems plausible to answer "high", "high" and "high". Specifically, we might be tempted to go for the following constraints:

$$c(\bigcirc(blA|inA)) > .5$$

$$c(\bigcirc(blB|inB)) > .5$$

$$c(\bigcirc(blN)) > .5$$

However, these constraints are probabilistically incoherent.[28] After all, the epistemic space is exhausted by inA and inB. In light of that, if blN has low probability conditional on each of inA and inB, it has to have low *unconditional* probability. This means that, if the premises of this argument are accepted, a theory of normative credence must be based on principles that are not classical principles of probability. This is not the

[28] To be more precise, they are incoherent relative to obvious background assumptions, such as the claim that blA, blB and blN are incompatible.

place to expand on these thoughts, except to express the hope for a clearer integration of a theory of normative uncertainty with deontic semantics.[29]

4 Puzzles of normality

Some of the classical puzzles of modal logic stem from the assumption that concepts of obligation are best modeled by normal modal operators. In particular, they stem from the assumption that *ought* expresses a concept of necessity. Philosophers of language have also reclaimed these puzzles as theoretical and empirical constraints in support of a variety of sophisticated theories. The two most important principles in this connection are:

INHERITANCE. $P \vdash Q \Rightarrow \bigcirc P \vdash \bigcirc Q$

AGGLOMERATION $\bigcirc(P)\ \&\ \bigcirc(Q) \vdash \bigcirc(P\ \&\ Q)$

These principles are elementary consequences of the principles that constitute normal modal systems—the rules of necessitation and substitution, together with axiom K for $\bigcirc$.

AXIOM K $\vdash (\bigcirc(P)\ \&\ \bigcirc(P \supset Q)) \supset \bigcirc(Q)$

We have already encountered an argument that implicitly targets normality. Specifically, in discussing Chisholm's paradox, I have taken note of Arregui's counterexample against DEONTIC DETACHMENT. This was the argument from (13)-(14) to (15). Suppose now that DD is best formalized as follows (for a conditional connective $\rightarrow$ validating *modus ponens*):

REFORMULATED DD $\bigcirc(P), \bigcirc(P \rightarrow Q) \vdash \bigcirc(Q)$

The combination of AGGLOMERATION and INHERITANCE entails this.[30] So any attempt at a counterexample to REFORMULATED DD must immediately be a counterexample to one of the two principles of normality.[31]

The literature also features more direct attacks to normality. Those attacks are what we focus on in the next two sections.

[29] For some related development of a non-factual theory of graded modal judgment (though not one that is especially focused on the case of deontic modality), see Charlow (forthcoming).

[30] By AGGLOMERATION, $\bigcirc(P), \bigcirc(P \rightarrow Q)$ entails $\bigcirc(P\ \&\ P \rightarrow Q)$. Because we choose $\rightarrow$ so that $P\ \&\ P \rightarrow Q \vdash Q$, INHERITANCE and $\bigcirc(P\ \&\ P \rightarrow Q)$ yield $\bigcirc(Q)$.

[31] There is room to escape this consequence if we refuse to formulate deontic detachment as in the above, and instead go in for NARROW DEONTIC DETACHMENT from Section 3.2.

4.1 Inheritance

Broadly speaking, there are two kinds of attempted counterexamples against INHERITANCE. The first consists of potential counterexamples against the entailment from $\bigcirc(P)$ to $\bigcirc(P \vee Q)$. The second consists of potential counterexamples against the entailment from $\bigcirc(P \ \& \ Q)$ to $\bigcirc(P)$. I will present them separately and then present the main conservative responses to them.

Disjunction inferences

Start with the disjunction side. The motivating observation is that Ross (1941)'s puzzle about imperatives extends to deontic modals. In his classic paper, Ross noted that imperatives like *mail the letter* do not seem to entail imperatives like *mail the letter or burn it*. Analogously, we might worry that (21) does not seem to entail (22)

(21) You should mail the letter.

(22) You should mail the letter or burn it.

The latter invites the inference that you may burn the letter, while the former does not. It is tempting to take this effect to be pragmatic and many theorists have endorsed this approach (Föllesdal and Hilpinen, 1971; Wedgwood, 2006). For example, Wedgwood suggests accounting for the badness of (22) in terms of a Gricean quantity implicature.

> There is an obvious Gricean explanation for why (22) seems an odd thing to say. It is much less informative than something else one might say–namely (21). Asserting the weaker claim would tend to be a useful contribution to a conversation only if one was not in a position to assert the stronger claim.

However, (Cariani, 2013a, §. 6) argues that it is implausible to take it as an pragmatic implicature, since it does not seem to behave as one.[32]

As an alternative, Cariani (2013a) offers a framework that could deal with this kind of phenomenon semantically. The design challenge that

[32] The dialectic interacts in significant ways with developments in the theory of scalar implicatures that are too rich to detail here. A popular approach, described for instance in Chierchia *et al.* (2008), treats scalar implicatures as entirely grounded in syntactic and semantic facts. The view in Cariani (2013a) is that if it turns out that there is a viable account of Ross-type phenomena within this framework, that would still be a win for the non-pragmatic camp, even if the theory turned out to have a different shape than initially anticipated.

needs to be met by any semantic framework of this kind is to explain how to account for the apparent falsity of (22) without landing on the evidently false hypothesis that $\bigcirc(\mathsf{P})$ can only be true if all the relevant P worlds are permissible. That hypothesis would lead to a vast proliferation of false *ought* judgments: consider, for instance, a true-sounding *ought* claim like *you ought to take good care of yourself*. There are plenty of overall impermissible ways taking good care of oneself, but they do not undermine the truth of the *ought*-claim.

Cariani (2013a) proposes that this be viewed as a form of 'coarseness' of the semantics. This is handled within a formalism that, taking a lead from Yalcin (2011), is referred to as *resolution semantics*. The key idea of that formalism is that the deontic orderings that the semantics accesses are not orderings of worlds but orderings of alternatives — where alternatives are coarser objects than worlds.[33]

The discussion of Cariani (2013a) fails to identify the complete logic of the resolution semantics for $\bigcirc$. This logical project has been taken up and completed by Van De Putte in a striking contribution (Van De Putte, 2018). Van De Putte shows that the non-normal obligation operator in Cariani (2013a) can be decomposed in three normal operators, each of which has a well-understood logic. In addition to the intrinsic interest of Van De Putte's results, his work stands as a model of what a two-way interaction between logic and philosophy of language might look like.

Conjunction inferences

As for conjunction elimination, the central puzzle case is Frank Jackson's Professor Procrastinate case Jackson (1985), Jackson and Pargetter (1986).[34] In Jackson's story, Prof. Procrastinate is asked to write a book review on a subject on which he is the foremost expert. However, Procrastinate's disposition is such that if he accepts the commitment to write the review, he won't write it. Jackson's judgments in this case are that *Procrastinate ought to accept the commitment and write the review* but *Procrastinate ought to accept* is false (because if he accepts he won't write).

[33]Some of my later work on information sensitivity Cariani *et al.* (2013), Cariani (2013b), Cariani (2016b), seeks to show that this idea has valuable applications even if one *accepts* INHERITANCE. Indeed, I maintain that this is critical to get a proper treatment of information-sensitivity. In addition, there are other applications of these contrastive ideas. For example, Snedegar (2017), mounts an impressive case that contrastivism helps solve some puzzles in the logic and semantics of reasons-claims.

[34]von Fintel (2012) notes that these cases have antecedents in a very similar case by (Kamp, 1973, pp.59-60) and are importantly like some cases that have been described in the literature on desire ascriptions.

Jackson thought that the intuitions surrounding the Procrastinate case are evidence for an *actualist* semantics, according to which $\bigcirc(\mathsf{P})$ is true at w if the value of the closest P-world to w exceeds the value of the closest $\neg\mathsf{P}$-world. This sorts of semantics can be viewed as an ancestor to Arregui's proposal. What they share is the idea that there is a counterfactual element in deontic modality.

Cariani (2009, 2013a) points out that the counterfactual element in the Procrastinate case is not necessary to get the relevant intuitions going. More specifically, we don't have to stipulate that Procrastinate *won't* write the review. All that is necessary is the stipulation that it is *very unlikely* that he will. Partly with this kind of puzzle in mind, Cariani (2009) and, much more extensively, Lassiter (2011, 2017) developed decision-theoretic accounts of the meaning of *ought*.[35] According to these decision-theoretic approaches, Procrastinate ought to accept and write the review because the expected value of accepting and writing is higher than the appropriate threshold. (The threshold itself might be understood contrastively as the expected value of the salient alternatives). By contrast, merely accepting is low in expected value, given the low probability of the good outcome if Procrastinate accepts.

That said, rejecting INHERITANCE in Procrastinate cases—even these probabilized ones—is not the sole province of expected utility accounts. Both the resolution semantics of Cariani (2013b) and Arregui's semantics for Chisholm's paradox make the same prediction about these cases, without appealing to expected utilities.

Responses and Arguments in Favor of Monotonicity

There are important criticisms for all of these moves by defenders of INHERITANCE. von Fintel (2012) objects that the disjunction in Ross's Puzzle is free choice disjunction[36] and that the intuitions in Procrastinate cases can be addressed as involving context shift, and specifically as

[35] For an initial presentation of a decision-theoretic semantic within the deontic logic literature, see Goble (1993) (though Goble does not apply the semantics to Procrastinate cases). For a very different kind of broadly decision-theoretic account—one which actually validates INHERITANCE—see Wedgwood (2016). For some reasons to avoid this decision-theoretic semantics, see the literature on semantic neutrality, including Cariani (2014, 2016a); Carr (2015); Charlow (2016, 2018).

[36] As I have noted, von Fintel suggests treating Ross's puzzle as free choice. This possibility was anticipated in (Cariani, 2013a, p. 551). The response provided there was that free choice arises equally with deontic and non-deontic interpretations of modals, but Ross-type phenomena don't. Intuitions about failures of INHERITANCE are much weaker for epistemic *must*, than they are for even deontic *must*. In fact, I suspect those intuitions can entirely be accounted for in pragmatic terms.

involving expansions of domains of quantification as more possibilities are made salient. We reject *Procrastinate ought to accept* when the possibility of writing is not salient. We accept *Procrastinate ought to accept and write* because that ideal possibility needs to be salient for this to be even evaluable. (Bronfman and Dowell 2018 also adduce similar context-shift considerations.) Von Fintel notes one consideration that supports the context-shift move. It is extremely bad to conjoin the two Procrastinate sentences. For example, (23) feels like a contradiction.

(23) # Procrastinate ought to accept and write, but he ought not to accept.

Von Fintel is right that any viable theory founded on the rejection of INHERITANCE must account for why these sentences sound contradictory. Little has been said by critics of INHERITANCE, myself included, to account for such judgments in a systematic way.[37]

Recent work has also elaborated on the connection between Ross's puzzle and free choice phenomena, though perhaps not quite in the direction envisaged by von Fintel (2012).[38] In particular, substantial progress was made in work by Fusco (specifically Fusco (2015, 2014)) who provides an analysis of permission in terms of ratifiability and a two dimensional semantics for disjunction that allows a unified account of Ross-type phenomena and free choice permission. Here is a quick sketch of Fusco's semantics. Points of evaluations are triples $\langle s, y, x\rangle$ consisting

[37] Von Fintel also notes another debt that such theories incur. There is a rich linguistic literature on the status of NPI's (Negative Polarity Items). These are items like *any* that are only allowed in special environments. The standard view is that what makes an enviroment "special" in the relevant sense is that it is downward monotonic (Here $\Phi(__)$ is downward monotonic iff $\Phi(Q) \models \Phi(P)$ whenever P entails Q.) The problem is that *any* is licensed in *You don't have to bring any alcohol to the party*. This is captured by the INHERITANCE validating semantics because $\neg \bigcirc(\cdot)$ is downward entailing iff $\bigcirc(\cdot)$ is upward entailing—that is iff $\bigcirc$ satisfies INHERITANCE.

I concede to von Fintel that many of the INHERITANCE-rejecting proposals don't do a good job of tackling this issue, and I won't do much better here. With that said, I am more optimistic about the prospects of addressing this concern. For one thing *any* is licensed in environments that are clearly *not* downward monotonic. One example is *It's fifty percent likely that Mary didn't eat any of the cookies*. So what really seems to matter is that the environment be in some sense *locally* downward entailing. That requires getting clear about what it is to be "locally" downward entailing. Depending on how *that* story goes it seems possible to reconstruct the semantics of *ought* so that it combines a non-monotonic component and a monotonic one.

[38] I lack the space for a summary of the free choice literature here, which goes well beyond the modal case. But *that* is also relevant to the theme "Deontic Logic and Natural Language". So, it is worthwhile identifying some of the key references. Kamp (1973) is justly heralded as a classic. Among recent works that are of specific importance to the study of deontic concepts are: Aher (2012); Barker (2010); Starr (2016); Willer (2018).

of a state s, the actual world y (unshiftable by modal operators), and the world of evaluation x. First, define an auxiliary concept Alt_w:

$$Alt_w(\mathsf{P},\mathsf{Q}) = \{\mathsf{R} \in \{\mathsf{P},\mathsf{Q}\} \mid \mathsf{R} \text{ is true at } w\}$$

Informally, this returns whichever subset of $\{\mathsf{P},\mathsf{Q}\}$ contains all and only the sentences that are true at w. Letting R be a deontic accessibility relation, say:

$$[\![\mathsf{P} \text{ or } \mathsf{Q}]\!]^{s,y,x} = 1 \text{ iff } \exists \mathsf{R} \in Alt_y(\mathsf{P},\mathsf{Q}), [\![\mathsf{R}]\!]^{s,y,x} = 1$$

$$[\![\bigcirc(\mathsf{P})]\!]^{s,y,x} = 1 \text{ iff for all worlds } v \in s \text{ s.t. } xRv,\ [\![\mathsf{P}]\!]^{s,y,v} = 1$$

Informally, a disjunction is true (relative to s, y and x) if one of the alternatives drawn from P and Q that hold at y is true (again, relative to s, y, and x). The semantics for $\bigcirc$ is relatively standard (except for the rich points of evaluation).

Let us follow along with this semantics' evaluation of (22)—*you ought to mail the letter or burn it*.

$$[\![\bigcirc(m \text{ or } b)]\!]^{s,y,x} = 1 \text{ iff for all worlds } v \in s \text{ s.t. } xRv,\ [\![m \text{ or } b]\!]^{s,y,v} = 1$$

$$[\![m \text{ or } b]\!]^{s,y,v} = 1 \text{ iff } \exists \mathsf{R} \in Alt_y(m,b), [\![\mathsf{R}]\!]^{s,y,x} = 1$$

Suppose s contains worlds in which you burn the letter without mailing it (not that it would be helpful to mail a charred letter). Suppose further one such a world is actual. Then the disjunction m or b will collapse on b, and so $\bigcirc(m \text{ or } b)$ will collapse on $\bigcirc(b)$ which can easily be made false even if $\bigcirc(m)$ is true. The story can be completed with a "ratifiability"-based account of permission that also delivers an account of free choice.

Fusco's view shares with the accounts by Jackson and Arregui discussed above the idea that what ought to be the case might depend in part on what actually is the case. However, in Fusco's system this effects stems crucially from the semantics for disjunction, and it is not especially linked to the semantics for obligation. This means in particular that her semantics treats Ross type phenomena as fundamentally different from apparent failures of conjunction elimination inside $\bigcirc$. Furthermore, for that very reason, the contingency that Fusco recognizes in deontic claims is unlike the contingency we saw at work in Chisholm's paradox.

4.2 Agglomeration

Next up is the delicate matter of the agglomeration principle.

AGGLOMERATION $\bigcirc(P)\ \&\ \bigcirc(Q) \vdash \bigcirc(P\ \&\ Q)$

There are different putative counterexamples to AGGLOMERATION, and it seems likely that these ought to be evaluated independently.

Deontic conflict

The most famous examples involve deontic conflicts. Simple, illustratives examples of deontic conflict involve moral dilemmas (but note that there are many deontic conflicts that are not dilemmas). Maybe all of the following are true: (i) I ought to travel to a different country to perform my civic duties, (ii) I ought to stay home to attend to my ailing mother but (iii) it's not the case that I ought to do both. The study of logics for normative conflicts is a remarkably well developed area of investigation in deontic logic (see Goble (2013) for a survey).

My (possibly controversial) opinion is that the philosophy of language literature and linguistics hasn't produced much that is new on deontic conflict. Most typically, theorists recognize the need for an account of deontic conflict, but relegate that account to an additional module to be separately injected in one's semantic framework.[39] This is not to say that the issue gets no substantive discussion. Lassiter (2011) sought to reduce dilemmas to a more general pattern of non-agglomerating *oughts*. But, as Lassiter himself recognizes in later work (Lassiter, 2017, §8.11), this cannot be exactly right. There is something distinctive going on in the case of dilemmas—something that requires additional treatment even if, as Lassiter does, we reject AGGLOMERATION for independent reasons.

Indeed (Lassiter, 2017, §8.11) is one of the few exceptions to my claim that new ideas in deontic dilemmas from the linguistics side are few and far between. In that section of *Graded Modality*, Lassiter explores the idea that there might be connections between the structural features that give rise to deontic conflicts and the structure of multi-dimensional adjectives. You might think for instance that *clever* tracks many dimensions of cleverness, and similarly you might think that deontic words (or, better, priority modal expressions generally) track many dimensions of priority. This is an interesting idea that, though still in the early stages, is likely to draw attention and development in the future. [40]

[39] Versions of this strategy are gestured at in Cariani (2013a) and von Fintel (2012).

[40] There is interesting and related work on how to extract deontic domains from inconsistent premise sets, e.g. Silk (2017). There is also important work concerning how deontic conflicts might figure in a theory of *reasoning*: see Nair (2014, 2016).

Other challenges to AGGLOMERATION.

Some authors (Jackson, 1985; Finlay, 2014; Lassiter, 2011) propose accounts of deontic *ought* that involve more extensive violations of AGGLOMERATION. That is to say, violations of AGGLOMERATION that go beyond the isolated case of deontic conflict. For example, Jackson proposes this kind of example:

> Attila and Genghis are driving their chariots towards each other. If neither swerves, there will be a collision; if both swerve, there will be a worse collision [...]; but if one swerves and the other does not, there will be no collision. Moreover if one swerves, the other will not because neither wants a collision. Unfortunately, it is also true to an even greater extent that neither wants to be 'chicken'; as a result what actually happens is that neither swerves and there is a collision. It ought to be that Attila swerves, for then there would be no collision. [...] Equally it ought to be that Genghis swerves. But it ought not to be that both swerve, for then we get a worse collision. (Jackson, 1985, p.189)

(Cariani, 2016a, p.400-401) expresses skepticism about treating this as a general counterexample to AGGLOMERATION.[41] But let us concede Jackson's description of the data, for the sake of presentation. Note that the example, in its intended reading, does not involve multiple conflicting conflicts of value. Jackson's *ought*-claims are understood as relative to a world-evaluation method according to which the best possible outcome is just to avoid collisions, and in which the worse the collision the worse the world. Given all that, if this is a counterexample to AGGLOMERATION, then there are counterexamples to AGGLOMERATION that do not arise from deontic conflicts.

Corresponding to these putative violations are semantic theories that block these instances of AGGLOMERATION. For Jackson, α ought to fly is true at world w if the nearest world in which α flies is better than the nearest world in which α does not fly. Agglomeration fails when the nearest P & Q-world is bad, but the nearest P-world is good (because it

[41] I don't doubt that there are strong judgments in the direction of Jackson's intuition. In fact, the naked effect is strongly supported by an experimental study by (Lassiter, 2017, §8.7). I doubt that Jackson's chicken case is well understood as a case in which all the contextual parameters are held fixed throughout the evaluation of the argument. I have similar views about the scenario in Lassiter's experiment. For a fuller articulation of this kind of diagnosis, see Boylan (ms.).

is a $P \ \& \ \neg Q$-world) and the nearest Q-world is also good (because it is a $\neg P \ \& \ Q$-world). For Lassiter, α ought to fly is true relative to a probability p and utility u function if the expected utility of flying (calculated on the basis of p and u) exceeds the expected utility of the alternatives.

Cariani (2016a) argues that the path traced by these proposals ends up in a problematic place. In particular, nearly every implementation of this idea ends up violating something much more plausible than AGGLOMERATION — namely the inference:

WEAKENING. $\bigcirc(P), \bigcirc(Q) \vdash \bigcirc(P \vee Q)$.

The few combinations that avoid this pitfall end up validating AGGLOMERATION, either in general or in Jackson's alleged counterexample that was supposed to be a key motivating factor.

In response, Lassiter (2017) makes two interesting suggestions: one idea (§8.10) is to stipulate away the counter-models. The move is not entirely unprincipled. After all, Lassiter notes that the countermodels described in Cariani (2016a) are merely abstract: they show that it is possible for the semantics to invalidate weakening given certain parameter values, without describing cases that would instantiate that structure. So, perhaps there is room to rule them out by excluding the appropriate parameter values as inadmissible. If it is true that the abstract countermodels are hard to connect to intuitive scenarios, the theory incurs no *predictive* cost. The other approach Lassiter explores is to take WEAKENING as invalid, but recover it via scalar mechanisms. For elaboration of this alternative approach, see §8.14 of Lassiter (2017).

The last recent development I want to highlight here is Boylan's (Boylan, ms.) argument that *epistemic* but not *deontic* oughts fail to agglomerate. I won't present Boylan's counterexamples to epistemic *ought* AGGLOMERATION here, so as to not steal his thunder. But if he is right, and if we want to give a unified semantics for epistemic and deontic *ought*, we will need to reach for an abstract semantics that invalidates AGGLOMERATION and then recover deontic AGGLOMERATION *via* some special claim in the theory of flavors. Boylan's paper develops one way to do this.

5 Varieties of deontic strength

Another important research thread concerns the varieties of deontic strength. Theorists frequently associate English *ought* and *must* with distinct concepts of obligation. The *ought* concepts are generally taken to

be weaker than the related *must*-concepts. Let us start, then, by formulating that hypothesis officially:

STRENGTH ASYMMETRY. unembedded *must*-claims are logically stronger than their *ought*-counterparts

This difference is sometimes represented in terms of the claim that *must* is a 'strong necessity modal' while *ought* is a 'weak necessity modal'. Because this formulation ignores the views that deny that *ought* is any kind of necessity modal, it is better to start by formulating the idea in more general terms.

The STRENGTH ASYMMETRY thesis has been discussed in the deontic logic literature, well before its appearance in natural language semantics (Sloman, 1970; McNamara, 2010). However, linguistic work (largely spawned by Copley (2006) and von Fintel and Iatridou (2008)) has provided additional theoretical and empirical depth to this discussion.

5.1 Motivating Data

Below is a simple paradigm, drawing on von Fintel and Iatridou's examples:

(24) You ought to wash your hands but you don't have to.

(25) #You must wash your hands but you don't have to.

(26) #You must (/have to) wash your hands, but it's not the case that you ought.

The perceived inconsistency of (25) suggests that *must* and *have to* have similar levels of strength. By contrast, the consistency of (24) suggests that *ought* is weaker than both *must* and *have to*. A moment's reflection on similar data reveals that *should* patterns with *ought* — in the sense that replacing *ought* with *should* in (24) and (26) generates analogous judgments.

These observations invite the view that there are two levels of deontic force. As postulated in STRENGTH ASYMMETRY, the stronger level is occupied by *must* and *have to*, and the weaker level is occupied by *ought* and *should*.

Before moving to the central theoretical questions, let us dwell on one more empirical aspect of the asymmetry. In light of STRENGTH ASYMMETRY, we may consider how permission claims pattern. In particular, data

resembling the paradigm from von Fintel and Iatridou suggests that English *may* is not the dual of *ought*. Indeed, if *may* has to be the dual of something, *must* seem to be the right choice.

(27) You ought to wash your hands, but you may not.
[*ought*(P) & ¬*may*(P)]

(28) # You must (/have to) wash your hands, but you may not.
[*must*(P) & ¬*may*(P)]

These considerations open the interesting data question whether *ought* has a dual in English (and indeed whether in other languages, the appropriate translations of *ought* happen to have dual).

This question has received substantially less attention than the STRENGTH ASYMMETRY, but it was recently addressed by Beddor (2017). Beddor uses the term "faultlessness" to express the dual of the concept of weak necessity (weak necessity being what he takes *ought* and *should* to denote). The linguistic question is whether faultlessness is linguistically realized, in English or other languages. Beddor goes on to propose that faultlessness is expressed in English by the phrase *is justified*, as it occurs in sentences like

(29) Ada isn't justified in attempting a cartwheel.

Partial support for this idea seems to come from the fact that (29) sounds roughly equivalent to *Ada should not attempt a cartwheel*.

We could check this by inspecting consistency judgments. The following are two key test cases for Beddor's conjecture.

(30) #Ada should attempt a cartwheel but she isn't justified in doing so.

(31) #?Ada is justified in attempting a cartwheel but she should not.

My own judgment is that (30) is unfixably bad, while (31) might be repaired by having the modal pick on something other than the initial set of reasons that provide the justification. This seems to accord with Beddor's hypothesis. If Beddor's hypothesis is correct, we would have the surprising consequence that some grammatical modals of English have lexical modals as their duals.

5.2 Accounts of the STRENGTH ASYMMETRY.

The central theoretical question concerning STRENGTH ASYMMETRY is what is the correct semantic explanation of the phenomenon. Recent literature has explored a variety of approaches. We will consider two types of approaches, and then mention an approach that does not quite fit either type. These are the two types:

> TYPE 1. *must* and *ought* are analyzed as quantifiers over worlds but associated with different domain-generation rules.
>
> TYPE 2. *must* and *ought* are degree expressions that are sensitive to different thresholds.

Implementations of these ideas are typically elaborated within very specific semantic frameworks. However, the broad ideas are sufficiently modular to allow presentation independently of those specific implementations. I refer the reader to the primary sources referenced below for detailed implementation.

The domain-centric approaches (TYPE 1) all involve the classical assumption that *ought*, *must* and their cognates are necessity operators of some sort. In this context, it make senses to refer to the former as 'weak necessity operators' and to the latter as 'strong necessity operators'. The central tenet of TYPE 1 theories is that domains of the weak operators must be subsets of the domains of the strong ones.[42]

Orientation check: if $D \subseteq D^+$, and $\Box$ and $\Box^+$ quantify respectively over D and D^+, then $\Box^+$ will be at least as strong as $\Box$. Indeed, under basic assumptions, it will be strictly stronger whenever the inclusion between domains is itself proper.

As von Fintel and Iatridou put the point:

> Our conception of weak necessity then makes them universal/necessity modals just as much as strong necessity modals are. What makes them weaker semantically is that they have a smaller domain of quantification: strong necessity modals say that the prejacent is true in all of the favored worlds, while weak necessity modals say that the prejacent is true in all of the very best (by some additional measure) among the favored worlds. (von Fintel and Iatridou, 2008, p.119)

[42] More precisely, if the domains vary from world to world (as they do in the baseline Kratzerian theory) the domains of the weak operators *at each world of evaluation* w are required be subsets of the domains of the strong ones at w.

The domains of quantification are generated by two layers of ordering sources. The *primary ordering source* contains those considerations that are hard requirements, while the *secondary ordering source* contains those considerations that are, in some sense, "additional criteria" that are significant for ranking worlds but not strictly required.

It is sometimes objected that this view does not explain the contrast. What it does, the objection goes, is merely writing the contrast into the semantics. More specifically one may press questions such as:

- How do we understand the talk of "hard requirements"?
- And what makes a criterion "additional"?

One attempt to address this challenge is found in Chrisman (2012). Chrisman proposes that the key distinction is between the *requirements* and the *recommendations* of morality. Of course, since *ought* and *must* may be used well beyond the domain of moral discourse, this distinction would have to be suitably generalized to cover all priority modals. In particular, for every source of priorities, we should be able to distinguish between what it requires and what it recommends.

Another substantive attempt to articulate a distinction in this neighborhood is Rubinstein (2012). Rubinstein (2012) proposes that *must* and *ought* signal different levels of commitment to the priorities that ground them. In particular, strong necessity modals track those priorities that are commonly accepted as such by conversational participants, while weak ones additionally track those items that are being proposed as additions to a to-do list, but are generally held to be up-for-grabs.

Let us move on to TYPE 2 approaches— those that propose that deontic modals are threshold expressions. These approaches start off by rejecting the idea that *ought* and *must* are necessity operators of any kind. Instead, they lean on an extended analogy between these modals and the probability operator *likely*. We could model *likely*(P) as claiming that P has probability greater than some threshold (perhaps .5). Similarly, we could model *ought*(P) as saying that P has a degree of 'oughtiness' greater than some threshold. Needless to say, an important sub-task consists in explaining what these "degrees of oughtiness" are. According to these views, we might model the difference between *ought* and *must* roughly on a treatment of threshold adjectives that are on the same scale but point to different threshold points (say *hot* vs. *warm*). For example, Lassiter (Lassiter (2011), Lassiter (2017)) proposes that both *ought* and *must* track expected values, but also that *must* demands higher thresholds than *ought*.

An alternative proposal in this family is found in Finlay (2016). For him, *ought* P means that P is more likely, conditional on the agent's contextually salient ends, than the alternatives. By contrast, *must* P means that P is certain, conditional again on the agent's contextually salient ends being realized. The asymmetric entailment is then recaptured by the simple fact that the truth-condition for *must* requires the truth-condition for *ought* to be satisfied.

These proposals generally do well at accounting for the core data supporting STRENGTH ASYMMETRY. (Some, but not all, TYPE 2 approaches have trouble with the permission data in (28), depending on how permission is analyzed.) Insofar as they run into trouble, it is because they run afoul of some other desiderata (see for example the discussion of AGGLOMERATION violation in Section 4.2 above).

The last account I want to highlight does not neatly fit either the TYPE 1 or the TYPE 2 mold. According to Silk (forthcoming), the fundamental difference between *must* and *ought* is that the former, but not the latter, requires that its prejacent be a necessity with respect to the actual priorities. Informally, *must* P says that P is necessary in the *actual* world; *ought* P says that P is necessary in a range of worlds that are relevantly related to the actual world but needn't *be* the actual world. There are two things that make this theory interesting: for one thing, it features a new way of capturing the "counterfactuality" of *ought*; for another, Silk is after the idea that there appear to be more levels of uncertainty about *ought* claims than there are about *must* claims. The most striking way in which Silk's account is at variance with the rest of the literature is that on his view *must* P does not entail *ought* P—so that *ought* and *must* are logically independent of each other. Silk gives some preliminary arguments (§4) that this is a feature and not a bug of his proposal (but the matter probably deserves greater scrutiny).

6 The grammar of action

It is common to assume that *ought*, *must* and *may* are propositional operators. The more carefully we integrate our analysis with linguistics, the more it is worth scrutinizing this assumption. An important development in this direction concerns the proper modeling of the interaction between agency and obligation in the context of a linguistically plausible model of deontic modality.

So here is some old news. Consider:

(32) Gaia should dance with Iris.

This is widely believed to have multiple interpretations, even fixing a broadly deontic flavor. On one interpretation, it identifies the ideal state by whatever salient criteria. In particular, it claims that the ideal state is one in which Gaia and Iris dance together. On another interpretation, it conveys that same content *plus* something extra—something to the effect that it is *up to* Gaia to make sure that she dances with Iris.

To elicit the first interpretation, consider a context in which the organizers of a ball are choosing which people are to dance with each other. It may well be best for the organizers that Gaia and Iris dance together, but it doesn't behoove Gaia to bring this about. To elicit the second interpretation, consider instead a context in which Gaia herself is choosing a dance partner. It is standard to label these the *ought-to-be* interpretation and the *ought-to-do* interpretation Feldman (1986). Though it is standard, however, this is not the only way of conceptualizing the dichotomy. Some authors distinguish between *agentive* interpretations and *non-agentive* interpretations of modals. Others contrast *deliberative* interpretations with *non-deliberative*, or *evaluative*, ones. The plurality of terminological choices suggests that it is probably a mistake to think of this as *one* dichotomy.

6.1 Logical approaches

There are many ways to design a formal system that is capable of generating a distinction like the one we are after. An unimaginative approach might stary by characterizing two sentential operators $\bigcirc_{be}$ and $\bigcirc_{do}$. Immediately, we might worry that this treatment does not attempt to explain the plurality of interpretations of (32). At best, the approach reflects some prior understanding of what the distinction might be.

What weight to give to this worry depends on what kind of application we have in mind. There are applications for which the multiple operator approach is entirely unobjectionable. After all, logic alone does not demand explanations of empirical phenomena. Moreover, a theorist might just be interested in questions like: what is the logic of *ought-to-be* operators?

Having said that, one doesn't need to be deeply invested with linguistics to ask more than what the two-operators approach can provide. For instance, one might want to design a system that can capture, without simply stipulating them, the logical interactions between deontic claims and sentences ascribing agency. This perspective is found in the familiar *stit* framework, as developed for instance in Horty and Belnap (1995);

Belnap *et al.* (2001); Horty (2001). In this framework, we add a family of sentential operators $stit_\alpha(\cdot)$ with the intended interpretation that agent α brings about the state of affairs that is described by the input proposition. So, $stit_\alpha(rain)$ might be a formalization of:

(33) α sees to it that it rains.

I will not develop a full *stit* framework since it is likely extremely familiar to the readers of this handbook. The key idea for our purposes is that *stit* framework allows us to capture the distinction between *ought-to-be* and *ought-to-do* in terms of a single obligation operator $\bigcirc$. This could be achieved by feeding $\bigcirc$ inputs that differ in their agentivity. So, $\bigcirc$(Gaia dance with Iris) symbolizes that it ought to be the case that Gaia dances with Iris. By contrast, $\bigcirc(stit_{Gaia}$(Gaia dance with Iris)) symbolizes that Gaia ought to see to it that she dances with Iris.

6.2 Linguistic concerns

Once again, this is unobjectionable logical development. But, for some applications, we need a more direct account of the connection between a formalism for *oughts* and agency and natural language.

In an influential paper Schroeder (2010) has charted a set of challenges for the *stit* analysis interpreted as a piece of philosophy of language.[43] As I interpret him, Schroeder argues that there are two linguistic constraints that are violated by the *stit* analysis. For future reference, let me give these constraint some names:

ADHERENCE. An *adherent* theory must predict, without overgenerating, which interpretations of any deontic sentence is available in any given context.

FAITHFULNESS. A *faithful* theory must respect, and ideally *account for*, relevant generalizations emerging from natural language syntax and semantics.

Schroeder argues that versions of the *stit* view generally violate both these desiderata.

The *stit* view fails to be adherent predicting agentive interpretations that are, in fact, not available. To see the force of the argument, we must first expand the *stit* view. Of course, we never say things like,

[43] In fairness this is not the spirit in which the analysis is typically proposed.

(34) It ought to be that Lisa sees to it that she runs.

We more often say things like:

(35) Lisa ought to run.

For the *stit* view to account for the agentive reading of (35), it needs to take on two additional claims. First, that a *stit*-operator is present but unpronounced in (35). Second, that the other implicit arguments that are required to fill out the logical form appropriately can also be reconstructed, even if they are not explicitly provided. In other words, (35) needs a logical structure roughly like (38). (Note: in (38) I use small caps for unpronounced material and parentheses to mark scope)

(36) $\bigcirc$ ($\textsc{stit}_{\textsc{lisa}}$ (Lisa run))

Finally, one must hypothesize that the new, unpronounced occurrence of LISA is somehow dependent on the overt subject of (35). This avoids the potential worry that (35) might end up being assigned the meaning that it is up to Simone to see to it that Lisa runs.[44] Call this expanded version of the *stit view* "the agency-in-the-prejacent" view" (AIP for short).

The problem Schroeder identifies is that AIP appears to overgenerate. Imagine a conversation in which two people emphatize with a friend, Luckless Larry, who has a had to endure a remarkable series of misfortunes. Suppose that one of them says:

(37) Larry ought to win the lottery.

According to Schroeder, (37) does not have an agentive interpretation—not, at least in contexts in which it is common ground that Larry has no ability to affect the lottery.[45] But it is unclear how AIP avoids this consequence. After all, the same mechanisms that generate (38) should generate a logical form roughly like:

(38) $\bigcirc$ ($\textsc{stit}_{\textsc{larry}}$ (Larry wins the lottery))

And that logical form should be assigned an agentive reading.

[44]More specifically, the worry is that it might otherwise get the logical structure $\bigcirc$ ($\textsc{stit}_{\textsc{simone}}$ (Lisa run)).

[45]Bronfman and Dowell (2018) push back on these judgments. According to them, the interpretation is not ruled out by the grammar, and instead it is simply made far fetched by contextual knowledge.

In addition—and perhaps more worryingly—Schroeder argues that AIP also violates FAITHFULNESS. Let us go back to the transition from (35) to (38). What must justify this transition is that *ought* is assimilated to a *raising* verb.[46] Without getting into the details of the syntax of raising constructions, the claim is that the relation between (38) and (34) parallels the relation between (39) and (40).

(39) Lisa appears to be tired.

(40) It appears that Lisa is tired.

In particular, the overt subject of (38) does not correspond to an argument of *appears*. Instead, it is in fact the subject of the "lower" verb, raised to the apparent position of subject of "higher" one. From the point of view of semantics, this means for instance that the semantic engine first evaluates the composition of *Lisa* and *be tired* and then applies the sentential operator *appears*.

Schroeder's point is that *ought* does not unambiguously meet the standard criteria for raising verbs. Contrast:

(41) Inez should examine June.

(42) June should be examined by Inez.

The hypothesis that *should* is a raising verb predicts that these are equivalent. This is why: suppose *should* were a raising verb, then the logical forms of (41)-(42) are roughly as in (43) and (44) respectively.

(43) ◯(Inez examine June)

(44) ◯(June be examined by Inez)

These can only have different truth-conditions if their prejacents have different semantic values. But their prejacents are related by passivization and it is a default assumption that passivization preserves semantic value—i.e., that sentences that are related by passivizations make the same contributions to the meanings of complex expressions that embed them. So, if *ought* is raising, (41)-(42) must have the same interpretations. a

Schroeder argues there seems to be a sense in which this is true, but that sense is not the one that is associated with agentive interpretations

[46]The category of *raising* verbs is from syntax where it is contrasted with the category of *control* verbs. Schroeder relies on textbook presentation from Radford (2004), which is as good as any reference.

of *ought*. Instead, that alternative sense seems closer to the *ought-to-be* interpretations. (To be exact, Schroeder avoids the *do/be* dichotomy and distinguishes between a *deliberative* and *evaluative* interpretation of *ought*.) One way to see this is that (41) and (42) in fact have *different* agentive interpretations. In its agentive interpretation, (41) demands that Inez see to it that she examine June. By contrast, (42) demands that June see to it that she is examined by Inez. So the agentive interpretations are not synonymous, and hence (41)-(42) cannot be raising verbs.

Schroeder's proposal is that *ought* is lexically ambiguous: one item, the "evaluative" *ought* is a sentential operator and classified as a raising verb; the other item, the "deliberative" *ought*, takes two separate arguments, an agent and an action.[47] In the possible worlds framework (which Schroeder does not adopt), the evaluative *ought* denotes a function from sets of possible worlds to truth-values.[48]

Schroeder's arguments have attracted a variety of responses. Chrisman (2012) argues that there are some diagnostics for control that are not satisfied by purportedly deliberative *oughts*. In particular, he focuses on the fact that typical control predicates can be nominalized. From *begin*, we can form *beginner*; from *want* we can (with some strain) form *wanter*; but it is ungrammatical to nominalize *ought*—there are no *oughters*. It is however unclear exactly how significant this observation is, since we already know that English *ought* is independently known to accept limited morphological combination (for example it lacks infinitive forms).

Other responses attempt to provide accounts of Schroeder's data that are nonetheless compatible with an ambiguity theory. Finlay and Snedegar (2014) propose a *contrastivist* account of the motivating data. For them, it is important to recognize that *ought*-sentences are relative to alternatives. Additionally, these alternatives are connected with focal stress. That is to say, that sentences like *Gaia ought to dance with Iris* suggests different things depending what speakers stress. Stressing *Gaia* suggests that the alternatives are possible dance partners for *Iris*. Stressing *dance* suggests that the alternatives are possible activities for Gaia. Stressing *Iris* suggests that the alternatives are possible dance partners for *Gaia*. The key claim is that in order to get a deliberative/agentive/ought-to-do reading, the

[47]There are quite a few antecedents for this kind of proposal in both deontic logic and lignuistics. See, among others, Thomason (1981a); Brennan (1993), and (Portner, 2009, §4.1.4 and §4.3.2).

[48]In the possible worlds framework (which Schroeder does not adopt), these can be assigned to types. Let t= the type of truth-values; e= the type of individuals and s=the type of possible worlds. Then the types of the two *oughts* are respectively $\langle\langle s,t\rangle,t\rangle$ and $\langle e,\langle e,\langle st\rangle\rangle,t\rangle$.

alternatives (generated by this or other similar means) must be in some sense "available" to the agent. This helps explain why we cannot get a deliberative interpretation of *Larry ought to win the lottery*. After all any such interpretation would have to include an unavailable option because winning the lottery is not (except in some exceptional cases) an available option.

Bronfman and Dowell (2018) defend a thoroughly contextualist approach by carefully working through Schroeder's arguments against ADHERENCE. Bronfman and Dowell show that the linguistic data underpinning Schroeder's argument are a little more sensitive to contextual variations than Schroeder initially suggested. Their theoretical proposal is to build on a relatively orthodox version of the baseline semantics. Their main contribution consists in articulating just how flexible the contextual apparatus is.[49]

In cases like (41) context supplies an ordering that tracks Inez's options and priorities; in cases like (42) context typically (but not invariably) supplies an ordering that tracks June's options and priorities. A natural challenge at this point is that this approach seems to stuff in a black box some systematic pattern that could be explained by the system. However, Bronfman and Dowell argue on the data side that the pattern is not *that* systematic after all

More recently, Daniel Skibra Skibra (unpublished) has leveraged a number of observations at the interface of syntax and semantics into an innovative and linguistically informed approach to the Schroeder's challenge. Like the other critics of Schroeder, Skibra maintains that *ought* is not lexically ambiguous between a deliberative and an evaluative interpretation. Unlike those critics, however, he believes that there is a difference in logical form between deliberative *sentences* and *agentive sentences*. In this respect, Skibra's view resembles the AIP account. However, he deviates from AIP because he thinks that the availability of agent arguments depends crucially on whether the *embedded verb* can take an agent argument. If Skibra is right, progress on this problem will require much deeper integration between the study of deontic modality and the theory

[49]To be clear, one could replicate their general strategy within many other semantic frameworks. Bronfman and Dowell present this as a defense of the baseline because given their work one *could* hold on to the baseline. Context sensitivity may well serve many masters.

of argument structure in syntax and semantics.

7 Conclusion

I have little in the way of general conclusions at the end of this general survey. We have considered some important ways in which attention to natural language phenomena has driven intellectual progress in deontic logic. The topic surveyed here are by no means the only ones deserving on emphasis. Among omissions I will emphasize some last themes:

- free choice permission.[50]
- work on the performativity of *must*.[51]
- work on the "Britney Spears" problem for deontic conditionals.[52]
- work on anankastic conditionals.[53]

The area remains active and vibrant.

[50] Kamp (1973); Fusco (2014); Starr (2016); Willer (2018).
[51] Ninan (2005); Mandelkern (forthcoming).
[52] Zvolenszky (2002, 2006); Cariani (2013a); Carr (2014).
[53] Sæbø (2001); Huitink (2008); von Fintel and Iatridou (2009); Condoravdi and Lauer (2016).

Bibliography

Aher, Martin (2012), 'Free choice in deontic inquisitive semantics (dis)', in 'Logic, Language and Meaning', 22–31, Springer.

Arregui, Ana (2010), 'Detaching *if*-clauses from *should*', *Natural Language Semantics*, 241–293.

Barker, Chris (2010), 'Free choice permission as resource-sensitive reasoning', *Semantics and Pragmatics*, 3, 10–1.

Beddor, Bob (2017), 'Justification as faultlessness', *Philosophical Studies*, 174, 901–926.

Belnap, Nuel, Perloff, Michael and Xu, Ming (2001), *Facing the Future*, Oxford University Press.

Bledin, Justin (2015), 'Modus ponens defended', *The Journal of Philosophy*, 112(2), 57–83.

Bledin, Justin (forthcoming), 'Fatalism and the logic of unconditionals', *Nous*.

Bonevac, Daniel (1998), 'Against conditional obligation', *Noûs*, 32(1), 37–53.

Boylan, David (ms.), 'Putting 'ought's together', Rutgers University.

Brennan, Virginia Mary (1993), *Root and epistemic modal auxiliary verbs*, Ph.D. thesis, University of Massachusetts.

Bronfman, Aaron and Dowell, Janice (2016), 'Contextualism about deontic modals', in Nate Charlow and Matthew Chrisman (eds.), 'Deontic Modality', 117–142, Oxford University Press.

Bronfman, Aaron and Dowell, Janice (2018), 'The language of 'ought' and reasons', in Daniel Star (ed.), 'Oxford Handbook of Reasons and Normativity', Oxford University Press.

Cariani, Fabrizio (2009), *The Semantics of 'ought' and the Unity of Modal Discourse.*, Ph.D. thesis, UC Berkeley.

Cariani, Fabrizio (2013a), '*Ought* and resolution semantics', Noûs, 47(3), 534–558.

Cariani, Fabrizio (2013b), 'Epistemic and deontic *Should*', *Thought*, 2(1), 73–84.

Cariani, Fabrizio (2014), 'Attitudes, deontics and semantic neutrality', *Pacific Philosophical Quarterly*.

Cariani, Fabrizio (2016a), 'Consequence and contrast in deontic semantics', *The Journal of Philosophy*, 113, 396–416.

Cariani, Fabrizio (2016b), 'Deontic modals and probabilities: One theory to rule them all?', in N. Charlow and M. Chrisman (eds.), 'Deontic Modals', 11–46, Oxford University Press.

Cariani, Fabrizio, Kaufmann, Magdalena and Kaufmann, Stefan (2013), 'Deliberative modality under epistemic uncertainty', *Linguistics and Philosophy*, 36, 225–259.

Carmo, José and Jones, Andrew JI (2002), 'Deontic logic and contrary-to-duties', in 'Handbook of philosophical logic', 265–343, Springer.

Carr, Jennifer (2014), 'The *If P, Ought P* problem', *Pacific Philosophical Quarterly*, 95(4), 555–583.

Carr, Jennifer (2015), 'Subjective ought', *Ergo*, 2(27), 678–710.

Charlow, Nate (2013), 'What we know and what to do', *Synthese*, 190, 2291–2323.

Charlow, Nate (2016), 'Decision theory: Yes! truth conditions: No!', in Nate Charlow and Matthew Chrisman (eds.), 'Deontic Modality', Oxford University Press.

Charlow, Nate (2018), 'Decision theoretic relativity in deontic modality', *Linguistics and Philosophy*, 41(3), 251–287.

Charlow, Nate (forthcoming), 'Grading modal judgment', *Mind*.

Charlow, Nate (unpublished), 'Modal triviality', .

Chierchia, Gennaro, Fox, Danny and Spector, Benjamin (2008), 'The grammatical view of scalar implicatures and the relationship between semantics and pragmatics', Harvard/MIT.

Chisholm, Roderick (1963), 'Contrary-to-duty imperatives and deontic logic', *Analysis*, 24, 33–36.

Chrisman, Matthew (2012), 'On the meaning of 'ought'', in Russ Shafer-Landau (ed.), 'Oxford Studies in Metaethics, vol. 7', 304, Oxford University Press.

Condoravdi, Cleo and Lauer, Sven (2016), 'Anankastic conditionals are just conditionals', *Semantics & Pragmatics*, 9.

Copley, Bridget (2006), 'What should *should* mean?', Unpublished manuscript, CNRS/University of Paris 8.

Dowell, Janice (2013), 'Flexible contextualism about deontic modals', *Inquiry*, 56(2-3), 149–178.

Dowell, Janice J. L. (2011), 'A flexible contextualist account of epistemic modals', *Philosophers' Imprint*, 11, 1–25.

Feldman, Fred (1986), *Doing the Best We Can: An Essay in Informal Deontic Logic*, Dordrecht Reidel.

Fine, Kit (unpublished), 'Chisholm's paradox and unconditional obligation', New York University.

Finlay, Stephen (2014), *Confusion of Tongues*, Oxford University Press.

Finlay, Stephen (2016), 'Ought: Out of order', in N. Charlow and M. Chrisman (eds.), 'Deontic Modals', 169–199, Oxford University Press.

Finlay, Steve and Snedegar, Justin (2014), 'One ought too many', *Philosophy and Phenomenological Research*, 84(1), 102–124.

von Fintel, Kai (1994), *Restrictions on Quantifier Domains.*, Ph.D. thesis, UMass, Amherst.

von Fintel, Kai (2012), 'The best we can (expect to) get? challenges to the classic semantics for deontic modals', presented at the 2012 Central APA, Chicago, IL, `http://mit.edu/fintel/fintel-2012-apa-ought.pdf`.

von Fintel, Kai and Iatridou, Sabine (2008), 'How to say *Ought* in foreign: the composition of weak necessity modals', in Jacqueline Guéron and Jacqueline Lecarme (eds.), 'Time and modality (Studies in Natural Language and Linguistic Theory 75)', 115–141, Springer.

von Fintel, Kai and Iatridou, Sabine (2009), 'What to do if you want to go to harlem?', MIT.

Föllesdal, Dagfin and Hilpinen, Risto (1971), 'Deontic logic: an introduction', in R. Hilpinen (ed.), 'Deontic Logic: Introductory and Systematic Readings', 1–35, Reidel.

Forrester, James (1984), 'Gentle murder,or the adverbial samaritan', *The Journal of Philosophy*, LXXXI(4), 193–197.

Fusco, Melissa (2014), 'Factoring disjunctions out of deontic modal puzzles', in Fabrizio Cariani, Davide Grossi, Joke Mehens and Xavier Parent (eds.), 'Deontic Logic and Normative Systems: DEON 2014', 95–107, Springer Verlag, Berlin.

Fusco, Melissa (2015), 'Deontic modality and the semantics of choice', *Philosopher's Imprint*, 1–27.

Gabbay, Dov (2012), 'Temporal deontic logic for the generalised chisholm set of contrary to duty obligations', in 'International Conference on Deontic Logic in Computer Science', 91–107, Springer.

Geurts, Bart (unpublished), 'On an ambiguity in quantified conditionals', Manuscript, University Njemegen.

Goble, Lou (1993), 'The logic of obligation, 'better' and 'worse'', *Philosophical Studies*, 70(2), 133–163.

Goble, Lou (1996), 'Utilitarian deontic logic', *Philosophical Studies*, 82, 317–357.

Goble, Lou (2013), 'Prima facie norms, normative conflicts, and dilemmas', in 'Hanbook of Deontic Logic and Normative Systems. Vol. 1', 241–351, College Publications.

Hacquard, Valentine (2006), *Aspects of modality*, Ph.D. thesis, MIT PhD.

Hacquard, Valentine (2010), 'On the event-relativity of modal auxiliaries', *Natural Language Semantics*, 18(1), 79–114.

Hacquard, Valentine (2013), 'The grammatical category of modality', in 'Proceedings of the 19th Amsterdam Colloquium', .

Hilpinen, Risto and McNamara, Paul (2013), 'Deontic logic: a historical survey and introduction', in 'Handbook of Deontic Logic and Normative Systems', 72–86, College Publications.

Holliday, Wesley and Icard, Thomas F. (2017), 'Indicative conditionals and dynamic epistemic logic', in 'Proceedings of TARK 2017', 337–351.

Holliday, Wesley and Icard, Thomas F. (2018), 'Axiomatization in the meaning sciences', in Derek Ball and Brian Rabern (eds.), 'The Science of Meaning', 73–97.

Horty, John F. (2001), *Agency and Deontic Logic*, Oxford University Press.

Horty, John F. and Belnap, Nuel (1995), 'The deliberative *Stit*: a study in the logic of action, omission, ability and obligation', *Journal of Philosophical Logic*, 24, 583–644.

Huitink, Janneke (2008), *Modals, conditionals and compositionality*, [Sl: sn].

Jackson, Frank (1985), 'On the semantics and logic of obligation', *Mind*, 94(374), 177–195.

Jackson, Frank (1991), 'Decision theoretic consequentialism and the nearest dearest objection', *Ethics*, 101(3), 461–482.

Jackson, Frank and Pargetter, Robert (1986), 'Oughts, options and actualism', *The Philosophical Review*, 95(2), 233–255.

Jerzak, Ethan (forthcoming), 'Two ways to want', *Journal of Philosophy*.

Jones, Andrew JI and Pörn, Ingmar (1985), 'Ideality, sub-ideality and deontic logic', *Synthese*, 65(2), 275–290.

Kamp, Hans (1973), 'Free choice permission', *Proceedings of the Aristotelian Society*, 74, 57–74.

Kaufmann, Stefan (2017), 'The limit assumption', *Semantics and Pragmatics*, 10(18), 1–27.

Kolodny, Niko and MacFarlane, John (2010), 'Ifs and oughts', *Journal of Philosophy*, 107(3), 115–143.

Kratzer, Angelika (1977), 'What 'must' and 'can' must and can mean', *Linguistics and Philosophy*, 1(3), 337–355.

Kratzer, Angelika (1981), 'The notional category of modality', in B. Partee and P. Portner (eds.), 'Formal Semantics: the Essential Readings', Blackwell.

Kratzer, Angelika (1991a), 'Conditionals', in A. von Stechow & D. Wunderlich (ed.), 'Semantics: An International Handbook of Contemporary Research', De Gruyter, from the *Semantics Archive*.

Kratzer, Angelika (1991b), 'Modality', in A. von Stechow & D. Wunderlich (ed.), 'Semantics: An International Handbook of Contemporary Research', De Gruyter.

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

Lassiter, Daniel (2011), *Measurement and Modality: the Scalar Basis of Modal Semantics.*, Ph.D. thesis, NYU.

Lassiter, Daniel (2016), 'Linguistic and philosophical considerations on bayesian semantics', in Nate Charlow and Matthew Chrisman (eds.), 'Deontic Modality', 82–116, Oxford University Press.

Lassiter, Daniel (2017), *Graded Modality*, Oxford University Press.

Levinson, Dmitri (2003), 'Probabilistic model-theoretic semantics for want', in 'Proceedings of SALT', volume 22, 222–39.

Lewis, David (1974), 'Semantic analyses for dyadic deontic logic', in 'Papers in Ethics and Social Philosophy', Cambridge University Press.

Loewer, Barry and Belzer, Marvin (1983), 'Dyadic deontic detachment', *Synthese*, 54(2), 295–318.

MacFarlane, John (2014), *Assessment Sensitivity*, Oxford University Press.

Mandelkern, Matthew (forthcoming), 'Practical Moore sentences', *Noûs*.

McNamara, Paul (2010), 'Deontic logic', in Edward N. Zalta (ed.), 'The Stanford Encyclopedia of Philosophy', Fall 2010 edition.

Nair, Shyam (2014), 'Consequences of reasoning with conflicting obligations', *Mind*, 123(491), 753–790.

Nair, Shyam (2016), 'Conflicting reasons, unconflicting "ought"s', *Philosophical Studies*, 173(3), 629–663.

Ninan, Dilip (2005), 'Two puzzles about deontic necessity', in Gajewski J., Hacquard V., Nickel B. and Yalcin S. (eds.), 'New Work on Modality', MIT Working Papers in Linguistics.

Parfit, Derek (unpublished), 'What we together do', Manuscript, Oxford University.

Portner, Paul (2009), *Modality*, Oxford University Press.

Prakken, Henry and Sergot, Marek (1996), 'Contrary-to-duty obligations', *Studia Logica*, 57(1), 91–115.

Prakken, Henry and Sergot, Marek (1997), 'Dyadic deontic logic and contrary-to-duty obligations', in 'Defeasible deontic logic', 223–262, Springer.

Radford, Andrew (2004), *Minimalist Syntax*, Cambridge Textbooks in Linguistics, Cambridge University Press.

Regan, Donald (1980), *Utilitarianism and Cooperation*, Oxford University Press.

Ross, Alf (1941), 'Imperatives and logic', *Theoria*, 7, 53–71.

Rubinstein, Aynat (2012), *Roots of modality*, Ph.D. thesis, UMass, Amherst.

Sæbø, Kjell Johan (2001), 'Necessary conditions in a natural language', *Audiatur vox sapientiae: A Festschrift for Arnim von Stechow*, 427–449.

Saint Croix, Catherine and Thomason, Richmond H. (2014), 'Chisholm's paradox and conditional oughts', in Fabrizio Cariani, Davide Grossi, Joke Mehens and Xavier Parent (eds.), 'Deontic Logic and Normative Systems: DEON 2014', 192–207, Springer Verlag, Berlin.

Schroeder, Mark (2010), 'Oughts, agents and actions', *The Philosophical Review*, 120(1), 1–41.

Silk, Alex (2014), 'Evidence sensitivity in deontic modals', *Journal of Philosophical Logic*, 43, 691–723.

Silk, Alex (2017), 'Modality, weights, and inconsistent premise sets', *Journal of Semantics*, 34(4), 683–707.

Silk, Alex (forthcoming), 'Weak and strong necessity modals', in Billy Dunaway and David Plunkett (eds.), 'Meaning, Decision, and Norms: Themes from the Work of Allan Gibbard', Maize Books.

Skibra, Daniel (unpublished), 'Ought and agency', Northwestern University.

Sloman, Aaron (1970), 'Ought and better', *Mind*, 79(315), 385–394.

Snedegar, Justin (2017), *Contrastive Reasons*, Oxford University Press.

Stalnaker, Robert. (1976), 'Indicative conditionals', *Philosophia*, 5(3), 269–286.

Starr, William (2016), 'Expressing permission', in 'Semantics and Linguistic Theory', volume 26, 325–349.

Szabó, Zoltán Gendler (2015), 'Major parts of speech', *Erkenntnis*, 80(1), 3–29.

Thomason, Richmond H. (1981a), 'Deontic logic and the role of freedom in moral deliberation', in R. Hilpinen (ed.), 'New Studies in Deontic Logic', 177–186, Reidel.

Thomason, Richmond H. (1981b), 'Deontic logic as founded on tense logic', in R. Hilpinen (ed.), 'New Studies in Deontic Logic', 165–176, Reidel.

Van De Putte, Frederik (2018), 'Coarse deontic logic', *Journal of Logic and Computation*, 1–32.

Veltman, Frank (1996), 'Defaults in update semantics', *Journal of Philosophical Logic*, 25, 221–261.

Wedgwood, Ralph (2006), 'The meaning of 'ought'', *Oxford Studies in Metaethics*, 1, 127–160.

Wedgwood, Ralph (2016), 'Subjective and objective ought', in N. Charlow and M. Chrisman (eds.), 'Deontic Modals', 143–168, Oxford University Press.

Willer, Malte (2012), 'A note on *Iffy* oughts', *Journal of Philosophy*, 109, 449–461.

Willer, Malte (2014), 'Dynamic thoughts on ifs and oughts', *Philosopher's Imprint*, 1–27.

Willer, Malte (2016), 'Dynamic foundations for deontic logic', in Nate Charlow and Matthew Chrisman (eds.), 'Deontic Modals', 324–354, Oxford University Press.

Willer, Malte (2018), 'Simplifying with free choice', *Topoi*, 37(3), 379–392.

Yalcin, Seth (2010), 'Probability operators', *Philosophy Compass*, 916–937.

Yalcin, Seth (2011), 'Nonfactualism about epistemic modality', in Egan A. and Weatherson B. (eds.), 'Epistemic modality', Oxford University Press.

Yalcin, Seth (2012), 'A counterexample to *Modus Tollens*', *Journal of Philosophical Logic*, 41, 1001–1024.

Zvolenszky, Zófia (2002), 'Is a possible worlds semantics of modality possible?', in B Jackson (ed.), 'Proceedings of SALT 12', 339–352.

Zvolenszky, Zófia (2006), 'A semantic constraint on the logic of modal conditionals', in Beáta Gyuris, László Kálmán, Chris Pi non and Károly Varasdi (eds.), 'Proceedings of the Ninth Symposium on Logic and Language (LoLa 9)', Révai Digital Press.