

Deontic modals are a form of normative language. They can be used express facts about deontic modality: to talk, not about about what is actually the case, but instead about what would be the case if the world accorded with some normative or evaluative standard. Various standards might be relevant: for example, moral norms, or requirements for achieving one's own private goals, or standards for complying with the mandates of good sportsmanship, or rules governing membership in JC's secret club.

Deontic modals are narrowly defined as a particular interpretation of modal auxiliaries: "ought", "may", "must", "can", "should", etc. But other sorts of expressions can express deontic modality:

- **verbs:** "have to", "need to", "be allowed to", "be supposed to", etc.
- **adjectives and adverbs:** "obligatory", "permissible", "impermissible", etc., as well as their adverbial forms
- **suffixes:** e.g. "desireable" to mean *worthy of desire*, rather than *possible to desire*
- **infinitival constructions:** "the thing to do", "you are to be home by 9:00", etc.

This chapter will focus strictly on deontic modals, but many of the conclusions generalize to other forms of normative and evaluative language (e.g., "good", "bad", "right", "wrong", "okay").

This chapter provides a selective survey of prominent theories of the semantics of deontic modals in logic and natural language. We focus on Kratzer's (1977; 1981; 1991) semantics and extensions to this analysis. Kratzer's semantics has been far and away the most influential theory of deontic modals, which provide a base case for the interpretation of normative language in general. Understanding the logic and truth-conditions of normative language is one of the core areas of metaethics. It informs our understanding of normative arguments and normative reasoning. As this chapter will emphasize, some forms of normative language don't allow for the inferences that classical logic trains philosophers to expect. Understanding what inferences are valid for normative language should impact our understanding of how we reason, and should reason, about the normative.

We will first look at how deontic modals are understood in the context of modal logic and natural language. Then we'll survey some recent debates and discoveries in the literature on deontic modals in natural language. We close with some considerations about the relevance of natural language to metaethics.

1 Logic and natural language

1.1 Deontic logic

In standard intensional logic, deontic modals are treated as a form of quantifier. They quantify over possible worlds. In a deontic context, 'Must α ' means, roughly, that α is true at *all* deontically ideal worlds. 'May α ' means, roughly, that α is true at *some* deontically ideal worlds. (I use '·' for quasi-quotation.) We can think of 'Must α ' as saying: *under any circumstances where things are perfect, α will be the case*. We can think of 'May α ' as saying: *it's compatible with things being perfect that α is the case*.

To be more precise, I'll briefly introduce a Kripke-style semantics for the deontic modals "may" and "must".

We start with a set of possible worlds W . Note that deontic logic is entirely neutral about the ontology of possible worlds. They could be understood as maximal sets of consistent sentences, or maximally specific properties, or concrete entities, etc. Talk of possible worlds might be realist or antirealist.

Standard deontic logic is concerned with how worlds are related to each other: which worlds are deontically ideal from the perspective of some other world. For example, if, in the actual world, a murder has been committed, the actual world is not deontically ideal. A deontically ideal world, from our perspective, might be one where the murder was not committed. It might be that there's one objectively privileged set of deontically ideal worlds. Or it might be that at different worlds, there are different norms or standards that determine different sets of worlds as deontic ideals. Logic is neutral on this question.

If a world w "sees" another world w' as deontically ideal— w' is ideal from the point of view of the standards at w —we'll say w' is "deontically accessible" from w . Call this relation of deontic accessibility D . If w' is deontically accessible from w , we'll say wDw' . For each w , there is a set of deontically

accessible worlds. We can call this set $D(w) = \{w' \in W : wDw'\}$.

Deontic modals are then understood as quantifiers over sets of deontically accessible worlds. We can distinguish deontic necessity modals from deontic possibility modals in terms of the type of quantification (universal or existential) they involve:

Deontic necessity modals: ‘must α ’, ‘should α ’, ‘ought α ’

Notation: $\Box\alpha$

‘ $\Box\alpha$ ’ is true at a world w iff α is true at all worlds in $D(w)$.

Deontic possibility modals: ‘may α ’, ‘can α ’

Notation: $\Diamond\alpha$

‘ $\Diamond\alpha$ ’ is true at a world w iff α is true at some world in $D(w)$.

Like “all” and “some”, \Box is traditionally assumed to be the dual of \Diamond . In other words, the two are interdefineable in the following way: $\Diamond\alpha$ can be defined as $\neg\Box\neg\alpha$, and $\Box\alpha$ can be defined as $\neg\Diamond\neg\alpha$. In English: α is permissible is true if and only if it’s not the case that not- α is required, and α is required if and only if it’s not the case that not- α is permissible.

The same kind of model can be used for other forms of modal logic, e.g. epistemic logic. Different logical subject matters, or modalities, can yield different constraints on the modal accessibility relation. For example, epistemic logic should be reflexive, in the sense that every world should be accessible to itself: for its accessibility relation, E , wEw for all w . This ensures that, where $\Box_s\alpha$ is read as *s knows that* α , $\Box_s\alpha$ is true only if α is true. But this is not true of deontic logic. People can sometimes fail to do what they ought to do, and so $\Box\alpha$ should not logically entail α . Deontic logic’s accessibility relation should not be reflexive.

Other constraints on a deontic accessibility relation are controversial. For example, it’s often suggested that deontic logic should be *serial*, such that every world has deontic access to some world: nothing is a deontic dead end. In other words, for all w there is some w' such that wDw' . This constraint ensures that for all propositions α , either $\Diamond\alpha$ or $\Diamond\neg\alpha$ (or both). Why does this matter? The seriality constraint rules out the possibility of *dilemmas*, where no options are permissible: both α and $\neg\alpha$ are

impermissible. Whether genuine dilemmas are possible is a substantive question: rather than being decided within the logic, we select a modal logic that conforms to first-order ethical considerations.

While this rudimentary deontic logic is adequate for many purposes, it remains highly idealized. It doesn't capture the ways in which uses of deontic modals in natural language are sensitive to features of context. It also doesn't capture the way in which deontic modals behave within different linguistic embeddings. (For a survey of deontic logics and discussion of its possible shortcomings, see (McNamara 2014).) So we turn to natural language considerations next.

1.2 Deontic modals in natural language

Like deontic logic, natural language semantics for deontic modals is of a piece with natural language semantics for modals in general. Epistemic, circumstantial, and other "flavors" of modality can be expressed with the some of the same vocabulary as deontic modality: *Jochen could drink whiskey* can be used to express the claim that Jochen was permitted to drink whiskey (deontic); the claim that, for all the speaker knows, Jochen might be a whiskey drinker (epistemic); or the claim that Jochen was capable of drinking whiskey (circumstantial).

There are any number of different flavors of modality expressible by modals. Even among deontic modals, different sets of norms or standards can be relevant for the interpretation of modalized assertions. For example, *Jochen could drink whiskey* might be true in a context where federal laws are the salient norms (he had just turned 21), but not in a context where other norms are salient (he was in a rehab program with strict rules about drug and alcohol consumption).

If we treated the multiple readings of modals as a form of ambiguity, we would have difficulties accounting for their relative cross-linguistic stability, as well as the commonalities in their embedding behavior. It makes better sense to treat the many readings of modals as a form of context-sensitivity, as Kratzer (1977) argued. This aligns with the semantics of other natural language quantifiers, the scope of which is also determined by context. If I say *All the students made it to the midterm*, context will determine whether I'm quantifying over all the students in the class, all the students in the university, or all the students in the universe. (In section 2.4, we will discuss which contexts are relevant for the assessment of deontically modalized sentences.)

The context-sensitivity of deontic modals explains why, in the context of federal law, *Jochen could drink whiskey* might be true, but (at the same time, of the same person) in the context of rehab regulations, the same sentence is not true. Federal law determines a different deontic accessibility relation from the relation determined by the regulations of the rehab clinic. Both sets of norms, we can suppose, determine a set of worlds consistent with their contents. But the worlds compatible with the federal laws include at least one where Jochen drinks whiskey, while the rehab regulations don't.

The picture, so far, is that deontically modalized sentences are evaluated relative to a context, which determines a set of deontically ideal worlds given the contextually salient norms or standards. The deontic modal acts as a quantifier over deontically ideal worlds.

Problem: suppose the following sentence is uttered in a context where the salient norms are the norms of morality:

- (1) You should rescue an abandoned dog (instead of buying a designer puppy).

Our current analysis of deontic modals entails that (1) is true if and only if all the morally ideal worlds are worlds in which you rescue an abandoned dog. But surely in the morally ideal worlds, there would be no abandoned dogs, since abandoning dogs is morally abhorrent. So this (highly plausible) sentence must be false! This is a simplification of Prior's (1958) Samaritan Paradox. (See also Åqvist 1967.)

One might reply: there's an easy solution. Context has made salient, here, *not* the morally ideal worlds, but the most *comparatively* morally ideal worlds that are consistent with the existence of abandoned dogs.

This explanation seems on the right track. Note that it appeals to two separate kinds of backgrounded information in a conversation: some facts about circumstances are taken for granted (in particular, the existence of abandoned dogs) and provide an absolute constraint on the domain of possible worlds that *should* quantifies over. (If there were a world with no abandoned dogs, it would be morally ideal, but in it, the addressee would not rescue a dog.) By contrast, the moral ideals do not place an absolute constraint: instead of looking to worlds that absolutely conform to moral norms, we look to those that conform as much as possible to moral norms consistent with the existence of abandoned dogs (and other circumstances taken as fixed). Moreover, this explanation suggests that the moral norms salient in the conversation are *comparative*: possible worlds can conform to the moral norms to

a greater or lesser extent. This suggests that these two forms of background information (what Kratzer (1981) calls “conversational backgrounds”) do not play formally equivalent roles.

These two observations lead to a semantics for deontic modals that incorporates separate parameters for conversational backgrounds: one for presupposed factual information, and one for an ordering over worlds in terms of some form of deontic ideality. This form of analysis is explored in (Hansson 1969; van Fraassen 1972; Lewis 1973). We will focus here on Kratzer’s (1981; 1991) semantics for modals in natural language, which incorporates the compositional context-sensitivity in (Kratzer 1977) and is explicitly unified with the restrictor analysis of conditionals.

On Kratzer’s analysis, the two conversational backgrounds that modals are sensitive to are the *modal base* and the *ordering source*. For deontic modals, the modal base is the set of circumstances that are taken as given (in our example, the existence of abandoned dogs, and probably other facts). The modal base is understood as a set of propositions. Kratzer treats propositions as sets of possible worlds. The modal base is thus a set of sets of possible worlds. Its intersection is the set of circumstantially possible worlds.

Ordering sources are also represented as sets of propositions, but their formal role is different: they are meant to generate a (partial) ordering over possibilities in terms of their ideality. An ordering can be projected from this set as follows: call our deontic ordering source g , and let $w \leq_g w'$ be interpreted as: *w is at least as ideal as w' according to g*. Then we can project the ordering \leq_g from g as follows:

$$w \leq_g w' \text{ iff } \{p \in g : w \in p\} \supseteq \{p \in g : w' \in p\}$$

In English: *w is at least as ideal as w' iff the set of ordering source propositions that w satisfies includes the set of ordering source propositions that w' satisfies.*

Illustrating with our dog example: suppose that the ordering is one of moral ideality. Simplifying morality a bit, our example could be represented with a comparative moral ideality relation as follows:

$$\begin{array}{c} \text{worlds where no one ever abandons a dog} \\ \wedge \\ \text{worlds where you rescue an abandoned dog} \end{array}$$

∧

worlds where there are abandoned dogs but you buy a designer puppy instead


How do we generate a domain for the modal to quantify over, given the modal base and the ordering source? First, we will assume that the domain of \leq_d must always have minimal (i.e., highest ranked) elements. (This is known as the Limit Assumption; see Lewis 1973 and Stalnaker 1984 for discussion.) Then the modal's domain is, roughly, the set of highest ranked worlds compatible with the modal base. More precisely: we can define the domain of a modal as a function of a world w , modal base f , and ordering source g :

$$\mathbf{domain}(w, f, g) = \{w' \in \cap f : \text{there is no } w' \text{ s.t. } w' \leq_g w \text{ and } w \not\leq_g w'\}$$

Then:

$\Box\alpha$ is true relative to $\langle w, f, g \rangle$ iff α is true at all w' in $\mathbf{domain}(w, f, g)$

$\Diamond\alpha$ is true relative to $\langle w, f, g \rangle$ iff α is true at some w' in $\mathbf{domain}(w, f, g)$



(A technical note: here we treat f and g as sets of propositions immediately generating a set of worlds and an ordering over worlds. But for reasons of compositionality, each of these conversational backgrounds will actually have to be functions from worlds to sets of propositions that generate (respectively) the set of worlds and the ordering over them. This allows that the same contextually salient circumstances and body of norms can generate different circumstantially possible worlds and different orderings at different worlds.)


So, in our example, the modal base eliminates worlds where no one ever abandons a dog. In context, we take it as given that abandoned dogs exist. Among the remaining worlds, worlds where you rescue an abandoned dog are better than worlds where you instead buy a designer puppy. So all the highest

ranked worlds in the intersection of the modal base are worlds where you rescue an abandoned dog. And so *You should rescue an abandoned dog* is true in this context.

One of the immediate benefits of this analysis is that it provides a solution to Kratzer’s (1991) version of the Samaritan Paradox. In Kratzer’s version, the paradox is to give a semantics for deontic modals that accommodates the following hypothesis: that (2a) and (2b) could be true in the very same context:

- (2) a. No one should abandon a dog.
- b. If there are abandoned dogs, you should rescue one of them rather than buying a designer puppy.

If deontic modals simply quantified over some immediately contextually specified domain—simply a salient set of deontically ideal worlds—then these two sentences couldn’t be true relative to the same context. For (2a) requires that the ideal worlds contain no abandoned dogs, and (2b) presupposes that at least one ideal world contains an abandoned dog. But Kratzer’s semantics makes sense of their consistency at a shared context. On her view, the function of the antecedent in (2b) is to restrict the modal base of the modal in the consequent to worlds where the antecedent is true. Where β is the conditional’s antecedent, the conditional can be represented with the following diagram:



So suppose our modal base’s intersection ($\cap f$) includes some worlds where there are no abandoned dogs and some worlds where there are. Then relative to the same conversational backgrounds, both (2a) and (2b) can be true: the highest g -ranked worlds in $\cap f$ are worlds where no one abandons dogs, and when $\cap f$ is restricted by the antecedent of (2b) (*there are abandoned dogs*), the highest g -ranked worlds are worlds where you rescue an abandoned dog rather than buying a designer puppy.

There are multiple linguistic phenomena associated with then name “The Samaritan Paradox,” and Prior’s original presentation is one that Kratzer’s semantics doesn’t immediately solve. The problem is

this: with both of the quantificational semantics we've considered, the following entailment holds:

- (1) You should rescue an abandoned dog (instead of buying a designer puppy).
- (3) Therefore, there should be abandoned dogs.

But this entailment is obviously a bad one; it is not inconsistent to endorse (1) and reject (3).

A simple (though not uncontroversial) solution to this puzzle is to place a *diversity condition* on the interpretation of modalized sentences:

Diversity: the intersection of the modal base where $\Box\alpha$ is evaluated must include α -worlds and *not*- α -worlds.

Diversity implies that (3) must be interpreted relative to a modal background that includes worlds where there are no abandoned dogs. But with the above assumptions about the ordering source, these worlds will be more highly ranked than worlds where the addressee rescues an abandoned dog. This means that (1) and (3) must be evaluated relative to different modal bases, and hence (1) cannot entail (3).

One potentially useful feature (or potentially problematic bug) in Kratzer's ordering sources is that they can be used to generate non-total orderings over worlds. This leaves room for incomparability between pairs of worlds: neither $w \leq w'$ nor $w' \leq w$. Consider, for example, a body of norms that included, among others, the following subset:

{you call your mother weekly, you refrain from speeding}

It might be that these two norms generate incomparability: if in w you satisfy only the first norm, and in w' you satisfy only the second norm, then maybe there's no fact of the matter about how w and w' compare in terms of deontically ideality. Since neither of these worlds satisfies a superset of the norms the other world satisfies, the partial order \leq projected from the ordering source g will not make a comparative ranking of the two.

Kratzer's semantics is the benchmark analysis of modals in natural language. Much of literature on deontic modals since has focused on elaborating generalizations of Kratzer's semantics to accommodate various puzzle cases. We'll survey some such cases now.

2 Puzzles in deontic semantics

2.1 Puzzles involving conditionals

The Gentle Murderer Paradox (Forrester 1984) consists of the following three premises, which are intuitively consistent:

- (4)
- a. You shouldn't murder anyone.
 - b. If you're going to murder someone, you should murder someone as painlessly as possible.
 - c. You're going to murder someone.

Classically, (4b) and (4c) entail, by modus ponens:

- (5) You should murder someone as painlessly as possible.

But (5) is, at least *prima facie*, inconsistent with (4a).

In the ethics literature, the most popular attempted solution to this problem aims to analyze the conditional (4b) in a way that makes modus ponens inapplicable in this case. On this analysis, the deontic modal *should* takes wide scope over the conditional: rather than having the logical form $\alpha \rightarrow \Box\beta$, (4b)'s logical form is $\Box(\alpha \rightarrow \beta)$. Because the main operator of the sentence is not a conditional, modus ponens doesn't apply. This wide-scoping proposal has been defended for a variety of similar problems in (Broome 1999, 2002, 2004; Dancy 2000; Darwall 1983, 2001). For a host of objections to this account, see, among others, (Kolodny & MacFarlane 2010) and (Silk 2014b). It is worth noting that Kratzer semantics does not have to be modified to make intuitive predictions about the Gentle Murderer case. Kratzer semantics doesn't validate modus ponens, as Charlow (2013) shows. For (4c) to be true at the world of evaluation, it is not necessary that it be incorporated into the circumstantial modal base relevant for the interpretation of (5). If it is not, then all the sentences in (4) are compatible.

Another puzzle generated by a tension between classical logic and deontic modals is the Miners Puzzle, first presented by Regan (1980) and later introduced into the semantics literature by Kolodny & MacFarlane (2010). The puzzle runs as follows:

Ten miners are trapped in one of two mineshafts, A or B, but you have no way to find out which; as far as you know, they're equally likely to be in either. Incoming floodwaters will

soon flood the shafts. You can block one or the other of the mineshafts using sandbags, but you cannot block both. If you block a mineshaft, all of the floodwaters will be diverted into the other, filling it completely and drowning any miners inside. On the other hand, if you block neither shaft, then both shafts will fill halfway with water, and only the one miner who is lowest in the mineshaft will be killed.

The sentences in (6) are intuitively true in this context:

- (6)
- a. You ought to block neither shaft.
 - b. If the miners are in shaft A, you ought to block shaft A.
 - c. If the miners are in shaft B, you ought to block shaft B.
 - d. The miners are either in shaft A or in shaft B.

The puzzle is that (6b), (6c), and (6d) classically entail (7) (by proof by cases):

- (7) Either you ought to block shaft A or you ought to block shaft B.

This example suggests that proof by cases might not be valid in natural language. Furthermore, it is impossible to accommodate within Kratzer's semantics without stipulating change of context between the sentences in (6).

Here's why: in Kratzer's semantics, the antecedent of a conditional can restrict the modal base, but it doesn't affect the ordering source. If (6a) is true, then the ordering source ranks blocking neither shaft more highly than blocking shaft A or blocking shaft B. But when the modal base is restricted to worlds where the miners are in shaft A (or in shaft B), the addressee still has the option of blocking neither shaft. And, since the antecedent doesn't affect the ordering source, it must still rank blocking neither shaft as the highest ranked option. So (6b) and (6c) are predicted to be false.

There are two strategies for addressing the Miners Puzzle. The conservative strategy explains away the Miners Puzzle as a case of shifting context; the inconsistency therefore turns on a kind of equivocation. The revisionary strategy explores how the Miners Puzzle sentences can be predicted within a single context, given a semantics that allows conditional antecedents to affect not just the modal base but also the ordering. Dowell (2012) and von Stechow (2012) provide conservative defenses of Kratzer semantics against the charge that the necessary changes in context are ad hoc. Revisionary generalizations of Kratzer semantics that accommodate the consistency of the sentences in (6) appear in (Kolodny

& MacFarlane 2010; Charlow 2013; Cariani, Kaufmann, & Kaufmann 2013; Silk 2014a; Carr forthcoming).

A final interesting puzzle for deontic modals and conditionals, which appears not to have a standard name, is what we will call the ‘If α , ought α ’ problem. The problem is that possible worlds semantics for modals, combined with various accounts of conditionals, predict that sentences of the form *if α , $\Box\alpha$* are logically valid, hence analytically true. Perhaps the earliest mention of this problem appears in (Spohn 1975); the problem is elaborated for Kratzer semantics in (Frank 1997; Zvolenszky 2002, 2006, 2007; Geurts 2004; Kratzer 2012; Carr 2014).

The problem is relatively straightforward: in Kratzer’s semantics for modals and conditionals, deontically modalized conditionals (of the form *if α , $\Box\beta$*) are true iff the highest ranked α -worlds in the intersection of the modal base are also β -worlds. Now, suppose β and α happen to be the same proposition. It is trivial that the highest ranked α -worlds in the intersection of the modal base are also α -worlds. So *if α , $\Box\alpha$* must be true.

But there are sentences of that form that are contingent, and moreover, intuitively false:

(8) If you cheat at Battleship, you ought to cheat at Battleship.

Now, an easy response is available. In section 1.2, we noted that Prior’s Samaritan Paradox appeared to motivate a diversity constraint on modal bases: $\Box\alpha$ can only be true if the intersection of the contextually determined modal base includes both α and $\neg\alpha$ worlds. But if the constraint applies to the modal base post-restriction when the modal is embedded in a conditional, this diversity constraint is violated in sentences like (8). And so with this generalization of the diversity constraint, we predict the falsity of (8). But we also predict that sentences of this form are necessarily false. And as Zvolenszky (2002) notes, some sentences of the form *if α , $\Box\alpha$* are intuitively true. For example, suppose Rita is a gentle spirit who seldom gets angry and only does so if she’s given very good reason. Then the following sentence is intuitively true:

(9) If Rita is angry, she ought to be angry.

While *if α , $\Box\alpha$* sentences are not trivially true, then, they are also not trivially false.

A better solution, mentioned in (Zvolenszky 2002; Kratzer 2012) and explored at length in (Geurts

2004), is to explore a kind of ambiguity generated by conditionals with quantifiers in their consequents. On Kratzer's semantics, bare conditionals (with no overt quantifier in the consequent) have an unpronounced quantifier (typically an epistemic modal) that is restricted by the antecedent. When a conditional has an overt quantifier, modal or otherwise, the conditional can be ambiguously read as also containing the unpronounced epistemic modal, generating different readings. On the doubly modalized reading, where the restrictor applies only to the epistemic modal, sentences of the form *if* α , $\Box\alpha$ are neither trivially true nor trivially false; and this accommodates the falsity of (8) and the truth of (9). (Zvolenszky 2002) argues that this account is inadequately predictive, however, and (Carr 2014) argues this account still faces problem cases.

2.2 Puzzles involving monotonicity

The next batch of puzzles for deontic modals involve monotonicity. Kratzer's quantificational semantics is upward monotonic. An operator ∇ is upward monotonic iff, if α implies β , then $\nabla\alpha$ implies $\nabla\beta$. A variety of examples suggest that natural language deontic modals might not be upward monotonic.

Jackson & Pargetter (1986) present a case of apparent failure of upward monotonicity involving conjunction:

Professor Procrastinate.

Professor Procrastinate receives an invitation to review a book...The best thing that can happen is that he says yes, and then writes the review when the book arrives. However, suppose it is further the case that were Procrastinate to say yes, he would not in fact get around to writing the review. Not because of incapacity or outside interference or anything like that, but because he would keep on putting the task off. (This has been known to happen.) Thus, although the best that can happen is for Procrastinate to say yes and then write, and he can do exactly this, what would in fact happen were he to say yes is that he would not write the review. Moreover, we may suppose, this latter is the worst that can happen. It would lead to the book not being reviewed at all, or at least to a review being seriously delayed. (235)

The following two sentences are judged to be true with respect to this case:

- (10) a. Professor Procrastinate should not accept the invitation.
 b. Professor Procrastinate should accept the invitation and write the review in a timely way.

So Professor Procrastinate is thought to show that $\Box(\alpha \wedge \beta)$ doesn't entail $\Box\alpha$.

Ross’s Puzzle (Ross 1944) involves an apparent failure of upward monotonicity involving disjunction: $\Box\alpha$ doesn’t entail $\Box(\alpha \vee \beta)$ (and similarly for \Diamond). Examples are easy to come by:

- (11) a. You should rescue an abandoned dog.
 b. So you should rescue an abandoned dog or burn down an animal shelter.

The inference from (11a) to (11b) is clearly not licensed.

One possible explanation for this phenomenon, discussed in (Cariani 2013) and defended in (von Fintel 2012), is that the inference is blocked because (11b) generates implications or implicatures that (11a) lacks—in particular:

- (12) a. You may rescue an abandoned dog.
 b. You may burn down an animal shelter.

This is a consequence of the so-called “free choice effect” (Kamp 1973):

- (13) a. $\Diamond(\alpha \vee \beta)$
 b. $\leadsto \Diamond\alpha \wedge \Diamond\beta$

Combining (13) with a seriality assumption, we validate the following inference:

- (14) a. $\Box(\alpha \vee \beta)$
 b. $\leadsto \Diamond\alpha \wedge \Diamond\beta$

(12b) is naturally judged false. Since (11b) licenses the inference to (12b) (by free choice), it must be false as well. But (11a) does not license the inference; so it is still judged to be true.

Now, it might be that (11b) is literally true; it is merely judged unassertable because its assertion would license free choice inferences. If these inferences are a purely pragmatic phenomenon (Hare 1967; Alonso-Ovalle 2006; Wedgwood 2006), rather than semantic entailments, then Ross’s puzzle generates no fundamental challenge to the compatibility of an upward monotonic logic for deontic modals. On the other hand, if free choice is semantic, then Ross’s puzzle suggests that deontic modals are not upward monotonic. (Cariani 2013; Lassiter 2011; Fusco 2015).

What about Professor Procrastinate? The case remains controversial. Cariani (2013) and Lassiter

(2011) defend the example as a failure of upward monotonicity, while others (e.g. von Fintel 1999) argue that the case involves a surreptitious context change. (10a), the claim that he should not accept the invitation, is evaluated relative to a context where it's treated as a necessary feature of the circumstances that Procrastinate will not write the review. But this proposition is no longer treated as necessary when we evaluate claim (10b), that Procrastinate should both accept the invitation and write the review. The question, once again, is whether to accept a conservative account that requires stipulating this context change, or instead a revisionary account that predicts the data without stipulations about context.

2.3 Strong and weak necessity modals

Throughout this chapter, we've assumed that different necessity modals—*should*, *ought*, *must*—were logically interchangeable, representable with the same logical symbol (\Box). As various authors have argued (e.g., Sloman 1970; Horn 1972; Ninan 2005; von Fintel & Iatridou 2008; Rubinstein 2012; Silk 2014a), however, expressions like *should*, *ought*, and *be supposed to* are importantly different from expressions like *must*, *need to*, *have to*, and *be required to*. The former group are logically weaker than the latter: (15a) is a perfectly acceptable utterance, while (15b) is hard to interpret as consistent: the assertion is infelicitous, though not in violation of grammar (conventionally marked with the “#” symbol).


- (15) a. Ian should rescue an abandoned dog, but he doesn't have to.
 b. #Ian has to rescue an abandoned dog, but it's not as if he should.

What is the difference between what's expressed with a so-called “strong necessity modal” like *must* and a so-called “weak necessity modal” like *should*?

Horn (1972) proposed a simple account: *must* α requires that α be true in all favored worlds, whereas *should* α requires only that α be true in most favored worlds. A benefit of this account is that it characterizes the relative weakness of *should*, between *may* and *must*, in the same terms that explain the logical weakness of *may* compared to *must*: in terms of their quantificational strength over a shared body of worlds. But the proposal also has shortcomings. It's not clear whether there's any sense to be made of counting possible worlds; even if there were, plausibly, the distinction between what one ought to do and what one must do is independent of the number of worlds. It might be that even among favored worlds, there could be many more ways for it to be the case that $\neg\alpha$ than for it to be the case

that α —but why should that make *should* α false?

von Fintel & Iatridou (2008) argue that *should* α says, not that most favored worlds are α -worlds, but rather that the best favored worlds are α worlds. This requires some ranking within the set of favored worlds. Modals must therefore be relativized to a second ordering. So, in our example, we suppose it is permissible for Ian not to adopt an abandoned dog: there are some favored worlds in which Ian does not adopt an abandoned dog. The ordering here could, for example be legal. But there also could be a second salient ordering—for example, in terms of maximizing overall utility—which, within the set of contextually favored worlds according to the first ordering, makes finer-grained distinctions.



Ninan (2005) describes another dimension of the strong/weak necessity modal distinction: that utterances of strong deontic necessity modals, unlike weak deontic necessity modals, seem incompatible with acknowledgement that the norms they express might not be obeyed.

- (16) a. Phaedra should feed the cat, but she won't.
 b. #Phaedra must feed the cat, but she won't.
 c. Phaedra should feed the cat, but she might not.
 d. #Phaedra must feed the cat, but she might not.

Ninan argues that the explanation for these data is that strong deontic necessity modals have a performative component in their conventional meaning: they are conventionally used not merely to report an obligation, but to impose one.

Given the distinctive uses of strong and weak necessity modals, questions arise about how the puzzle cases discussed in 2.1 and 2.2 apply when the type of necessity modal is toggled (see, e.g., Silk 2014a).

2.4 Contextualism and relativism

We now turn to the metasemantics of deontic modals. It's clear that these expressions, like other normative and evaluative vocabulary, are sensitive to features of context. But it is controversial which contexts they are sensitive to.

Contextualism about deontic modals is the view that utterances of deontic modals are sensitive to features of the context in which they are uttered. (For canonical defenses of contextualism about related expressions, though not specifically about deontic modals, see Kratzer 1981 and DeRose 1991.)

Contextualism: whether an utterance of $\Box\alpha$ at a context c is true depends on salient background information at c .

Relativism is the view that utterances of deontic modals are also sensitive to features of the context in which the utterances are assessed, which can be different from the context in which they are uttered. There are different ways this view can be formulated:

Truth relativism: whether an utterance of $\Box\alpha$ at a context c is true as assessed from a context c' depends on salient background information at c' .

Content relativism: what proposition an utterance of $\Box\alpha$ at a context c expresses, as assessed from a context c' , depends on salient background information at c' .

To see how these three views come apart, consider the following simplified example:

Disagreement.

Walker and Elle are considering whether how their friend should get to a different part of Chicago: on foot or on the L. Walker thinks the beauty of the walk makes up for the slower commute, while Elle thinks the opposite: the speed of taking the L compensates for the worse views. Walker asserts, in context c :

(17) He should walk.

Elle assesses (17) as false in context c' .

At which context(s) is (17) true? At which context, if any, is (17) false?

According to *contextualism*, the context of assessment can only be the context of utterance, c . At c , where the salient standards are those endorsed by Walker, (17) expresses, roughly, the proposition that according to Walker's standards, the friend should walk. This proposition is true.

According to *truth relativism*, (17) expresses the same proposition at both c and c' . But this proposition, and (17) itself, can only be assessed relative to contexts of assessment, which (at least in this case) determine a salient body of standards. The sentence and its corresponding proposition are true as assessed at c (relative to Walker's standards) but false as assessed at c' (relative to Elle's standards).

According to *content relativism*, (17) expresses different propositions as assessed at c and c' . Roughly, at c , (17) expresses the proposition that according to Walker's standards, the friend should walk. This proposition is true; hence at c , (17) is true. At c' , (17) expresses the proposition that according to Elle's standards, the friend should walk. This proposition is false; hence at c , (17) is false.

Truth relativism (with respect to a variety of expressions, including deontic modals) is canonically defended in (MacFarlane 2005, 2014) and (Egan 2007). Content relativism (about epistemic rather than deontic modals) is defended in (Egan et al. 2005) and (Weatherson 2009) and discussed in (Egan & Weatherson 2009), (MacFarlane 2014). For a more thorough discussion of arguments for and against contextualism, see chapter 6 of this volume. For arguments for and against different forms of relativism, see chapter 7 of this volume.

3 Deontic modals in ethics and metaethics

In this chapter, we have focused on the natural language logic and semantics of deontic modals. But one might worry that, for the purposes of ethical and metaethical theorizing, investigations of natural language are moot. For in ethical and metaethical theorizing, we regiment our use of deontic modals, so that, for example, the so-called objective *ought* always expresses non-information-sensitive, all-things-considered obligations, and the so-called subjective *ought* always expresses all-things-considered obligations that are relativized to the information of the agent(s) under discussion. For these purposes, one might think, the messiness of natural language can be idealized away, and the elegantly simple deontic logic and semantics described in section 1.1 will be adequate.

While we might allow that this regimented reading of deontic modals exists, if only within philosophy papers and philosophy classrooms, we should acknowledge that it is a philosopher's invention. We might reasonably be skeptical of the idea that we have intuitions about the truth values of claims involving a technical reading of deontic modals. Note, for example, that our reasoning and intuitions in ethics and metaethics often must be expressed using deontic modals under a variety of different embeddings. Can we really expect to have reliable intuitions about the truth conditions of the infinitely many complex sentences that our regimented deontic modals could appear in?

As a methodological point: native speaker intuitions are the primary data for facts about natural language; their correctness is the null hypothesis. Indeed, if linguistic facts are a matter of convention, then the aggregate of native speaker intuitions can't be wrong. The same methodological considerations don't apply to stipulated terms. We would need to stipulate the specific ways in which such terms are affected by embedding, rather than rely on intuition. Otherwise, our "intuitions" could (and, as a matter of sociological fact, probably do) reflect an unholy mixture of stipulation and intuition about natural language cognate terms.

Natural language is our primary mode of voicing our intuitions and expressing our reasoning in ethics and metaethics. And so we might believe we're using philosophically regimented deontic modals when we're not. This can lead to a variety of philosophical confusions: for example, misguided conclusions based on argument forms that are not valid within the logic of natural language deontic modals (like the Gentle Murderer), or illusions of paradox based on misguided assumptions about consistency and inconsistency (like the Miners' Puzzle).

Note that this argument does not necessarily extend to any use of regimented terms. For example, take Lewis's regimentation of the term "natural": the term applies to properties such that the "perfectly natural" properties are those that carve nature at the joints. The phrase "perfectly natural" is not context-sensitive; embedding it in different linguistic contexts will not affect its extension. So the above argument does not motivate research on the semantics of the natural language expressions "natural" or "perfectly natural." By contrast, modals are sensitive to a number of features of context.

Moreover, modals are standardly understood as logical operators. Logical operators are understood in terms of the inferences they warrant. For this reason, understanding their embedding behavior is

paramount for understanding their meaning. For cases like the Gentle Murderer or the Miners' Puzzle, even if we stipulate away the context sensitivity of deontic modals in our regimentation, we still turn to intuitions about consistency and entailment. But if consistency and entailment facts are not stipulated, then we risk confusing the embedding behavior and inferential role of the context-sensitive, natural language term with that of the stipulation. Indeed, whatever intuitions we might have about regimented readings of deontic modals, we should expect them to be systematically distorted by our intuitions about ordinary language.

It may be that some regimented use of deontic modals will succeed better than natural language at cutting nature at its ethical or logical joints. But we won't know unless we know how the regimented term could differ from its natural language counterpart. So it is worth the effort to get a clear grip on the logic and semantics of natural language normative expressions.

Acknowledgments

The research for this chapter has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP/2007-2013) / ERC Grant Agreement n. 312938. Many thanks to all with whom I've discussed these issues, and especially to Jamie Fritz, Tristram MacPherson, David Plunkett, and Paolo Santorio.

References

- Åqvist, Lennart (1967). "Good Samaritans, Contrary-to-Duty Imperatives, and Epistemic Obligations." *Noûs*, 1(4): pp. 361–379.
- Alonso-Ovalle, Luis (2006). *Disjunction in alternative semantics*. Ph.D. thesis, University of Massachusetts at Amherst.
- Broome, John (1999). "Normative Requirements." *Ratio*, 12(4): pp. 398–419.
- Broome, John (2002). "Practical Reasoning." In J. L. Bermúdez, and A. Millar (eds.) *Reason and Nature: Essays in the Theory of Rationality*, Oxford University Press, pp. 85–111.

- Broome, John (2004). "Reasons." In R. J. Wallace (ed.) *Reason and Value: Themes From the Moral Philosophy of Joseph Raz*, Oxford University Press, pp. 2004–28.
- Cariani, Fabrizio (2013). "Ought' and Resolution Semantics." *Noûs*, 47(3): pp. 534–558.
- Cariani, Fabrizio, Magdalena Kaufmann, and Stefan Kaufmann (2013). "Deliberative Modality Under Epistemic Uncertainty." *Linguistics and Philosophy*, 36(3): pp. 225–259.
- Carr, Jennifer (2014). "The 'If P, Ought P' Problem." *Pacific Philosophical Quarterly*, 95(4): pp. 555–583.
- Carr, Jennifer (forthcoming). "Subjective Ought." *Ergo*.
- Charlow, Nate (2013). "What We Know and What to Do." *Synthese*, 190(12): pp. 2291–2323.
- Dancy, Jonathan (2000). *Practical Reality*. Oxford University Press.
- Darwall, Stephen L. (1983). *Impartial Reason*. Cornell University Press.
- Darwall, Stephen L. (2001). "Because I Want It." *Social Philosophy and Policy*, 18(2): pp. 129–153.
- DeRose, Keith (1991). "Epistemic Possibilities." *Philosophical Review*, 100(4): pp. 581–605.
- Dowell, J. L. (2012). "Contextualist Solutions to Three Puzzles About Practical Conditionals." In *Oxford Studies in Metaethics, volume 7*, Oxford.
- Egan, Andy (2007). "Epistemic Modals, Relativism and Assertion." *Philosophical Studies*, 133(1): pp. 1–22.
- Egan, Andy, John Hawthorne, and Brian Weatherson (2005). "Epistemic modals in context." In G. Preyer, and P. Peter (eds.) *Contextualism in Philosophy*, Oxford: Oxford University Press, p. 131–168.
- Egan, Andy, and Brian Weatherson (2009). *Epistemic Modality*. Oxford University Press.
- Forrester, James William (1984). "Gentle Murder, or the Adverbial Samaritan." *Journal of Philosophy*, 81(4): pp. 193–197.
- Frank, Anette (1997). *Context Dependence in Modal Constructions*. Ph.D. thesis, Universität Stuttgart.

- Fusco, Melissa (2015). "Deontic Modality and the Semantics of Choice." *Philosophers' Imprint*, 15(28).
- Geurts, Bart (2004). "On an Ambiguity in Quantified Conditionals." Manuscript.
- Hansson, Bengt (1969). "An Analysis of Some Deontic Logics." *Noûs*, 3(4): pp. 373–398.
- Hare, R. M. (1967). "Some Alleged Differences Between Imperatives and Indicatives." *Mind*, 76(303): pp. 309–326.
- Horn, Laurence (1972). *On the Semantic Properties of Logical Operators in English*. Ph.D. thesis, University of California, Los Angeles.
- Jackson, Frank, and Robert Pargetter (1986). "Oughts, Options, and Actualism." *Philosophical Review*, 95(2): pp. 233–255.
- Kamp, Hans (1973). "Free Choice Permission." *Proceedings of the Aristotelian Society*, 74: pp. 57–74.
- Kolodny, Niko, and John MacFarlane (2010). "Ifs and Oughts." *The Journal of Philosophy*, 107(3).
- Kratzer, Angelika (1977). "What 'Must' and 'Can' Must and Can Mean." *Linguistics and Philosophy*, 1(3): pp. 337–355.
- Kratzer, Angelika (1981). "The Notional Category of Modality." In H.-J. Eikmeyer, and H. Rieser (eds.) *Words, Worlds, and Contexts: New Approaches in World Semantics*, Berlin: de Gruyter.
- Kratzer, Angelika (1991). "Modality." In A. von Stechow, and D. Wunderlich (eds.) *Semantics: An International Handbook of Contemporary Research*, Berlin: de Gruyter, pp. 639–650.
- Kratzer, Angelika (2012). *Modals and Conditionals*. Oxford: Oxford University Press.
- Lassiter, Daniel (2011). *Measurement and Modality: the Scalar Basis of Modal Semantics*. Ph.D. thesis, New York University.
- Lewis, David (1973). *Counterfactuals*. Oxford: Blackwell.
- MacFarlane, John (2005). "Making Sense of Relative Truth." *Proceedings of the Aristotelian Society*, 105(3): pp. 321–339.

- MacFarlane, John (2014). *Assessment Sensitivity: Relative Truth and its Applications*. OUP Oxford.
- McNamara, Paul (2014). "Deontic Logic." In E. Zalta (ed.) *Stanford Encyclopedia of Philosophy (Winter 2014 Edition)*. URL = <http://plato.stanford.edu/archives/win2014/entries/logic-deontic/>.
- Ninan, Dilip (2005). "Two Puzzles About Deontic Necessity." In J. Gajewski, V. Hacquard, B. Nickel, and S. Yalcin (eds.) *New Work on Modality, MIT Working Papers in Linguistics*, vol. 51.
- Prior, A. N. (1958). "Escapism: The Logical Basis of Ethics." In A. I. Melden (ed.) *Journal of Symbolic Logic*, University of Washington Press, 4, pp. 610–611.
- Regan, Donald (1980). *Utilitarianism and Cooperation*. New York: Oxford.
- Ross, Alf (1944). "Imperatives and Logic." *Philosophy of Science*, 11(1): pp. 30–46.
- Rubinstein, Aynat (2012). *Roots of Modality*. Ph.D. thesis, University of Massachusetts, Amherst.
- Silk, Alex (2014a). "Evidence Sensitivity in Weak Necessity Deontic Modals." *Journal of Philosophical Logic*, 43(4): pp. 691–723.
- Silk, Alex (2014b). "Why 'Ought' Detaches: Or, Why You Ought to Get with My Friends (If You Want to Be My Lover)." *Philosophers' Imprint*, 14(7).
- Sloman, Aaron (1970). "'Ought' and 'Better.'" *Mind*, 79(315): pp. 385–394.
- Spohn, Wolfgang (1975). "An Analysis of Hansson's Dyadic Deontic Logic." *Journal of Philosophical Logic*, 4(2): pp. 237–252.
- Stalnaker, Robert (1984). *Inquiry*. Cambridge, MA: MIT Press.
- van Fraassen, Bas (1972). "The Logic of Conditional Obligation." *Journal of Philosophical Logic*, 1(3/4): pp. 417–438.
- von Fintel, Kai (1999). "NPI Licensing, Strawson Entailment, and Context Dependency." *Journal of Semantics*, 16(2): pp. 97–148.

- von Fintel, Kai (2012). “The Best We Can (Expect to) Get? Challenges to the Classic Semantics for Deontic Modals.”
- von Fintel, Kai, and Sabine Iatridou (2008). “How to Say *Ought* in Foreign: the Composition of Weak Necessity Modals.” In *Time and Modality: Studies in Natural Language and Linguistic Theory*, vol. 75, pp. 115–141.
- Weatherson, Brian (2009). “Conditionals and Indexical Relativism.” *Synthese*, 166(2): pp. 333–357.
- Wedgwood, Ralph (2006). “The Meaning of the ‘Ought.’” In R. Shafer-Landau (ed.) *Oxford Studies in Metaethics: Volume 1*, Clarendon Press, vol. 1, pp. 127–160.
- Zvolenszky, Zsófia (2002). “Is a Possible-Worlds Semantics of Modality Possible? A Problem for Kratzer’s Semantics.” In *Proceedings of Semantics and Linguistic Theory 12*.
- Zvolenszky, Zsófia (2006). “A Semantic Constraint on the Logic of Modal Conditionals.” In B. Gyuris, L. Kálmán, C. Pinon, and K. Varasdi (eds.) *Proceedings of the Ninth Symposium on Logic and Language (LoLa 9)*, pp. 167–177.
- Zvolenszky, Zsófia (2007). *The Lost Pillar of Deontic Modality Naming with Necessity (Part of the Dissertation Portfolio Modality, Names and Descriptions)*. Ph.D. thesis, New York University.