Deontic Modals Without Decision Theory

Jennifer CARR — *Massachusetts Institute of Technology*

Abstract. The classic account of modals faces counterexamples. It makes the wrong predictions for a wide range of cases involving information-sensitive deontic modals. Some conservative amendments to the classic account have been proposed in order to avoid these problems. These accounts also face counterexamples. I argue that these counterexamples are manifestations of a deeper problem for the classic account and its recent variants: they all inadvertently build controversial normative assumptions into the semantics of modals. These normative assumptions come in the form of decision rules: they tell us how to go from some objective body of values to a verdict about what subjectively we ought to do, given our limited information. The fact that these decision rules are unattractive explains why many of the resulting predictions are judged false. I propose a generalization of the classic account. The view I defend introduces an additional parameter that is sensitive to norms of rational action under uncertainty. Instead of building these norms into the semantics, we should let them be determined by context.

1. Introduction

The classic account of modals (Kratzer 1981, Kratzer 1991) faces counterexamples. It makes the wrong predictions for a wide range of cases involving information-sensitive deontic modals. Some conservative amendments to the classic account have been proposed in order to avoid these problems (Cariani, Kaufmann, & Kaufmann 2011, Charlow 2012). These accounts also face counterexamples. I argue that these counterexamples are manifestations of a deeper problem for the classic account and its recent variants: they all inadvertently build controversial normative assumptions into the semantics of modals. These normative assumptions come in the form of decision rules: they tell us how to go from some objective body of values to a verdict about what subjectively we ought to do, given our limited information. The fact that these decision rules are unattractive explains why many of the resulting predictions are judged false. I propose a generalization of the classic account. The view I defend introduces an additional parameter that is sensitive to norms of rational action under uncertainty. Instead of building these norms into the semantics, we should let them be determined by context.

2. The classic account and a counterexample

2.1. Setting up the account

The near-orthodox account of conditionals in ordinary language is Kratzer's (1981, 1991) restrictor analysis. On the restrictor analysis of conditionals, all indicative conditionals contain a (sometimes covert) modal operator. The antecedent restricts the domain of that operator. Kratzer endorses a pairing of the restrictor analysis of conditionals with what von Fintel (2012) calls "the classic

account of modals."

The classic account uses two parameters, which Kratzer calls a "modal base" and an "ordering source." On Kratzer's interpretation of these parameters, the modal base is a set of propositions, the intersection of which is a set of (epistemically, circumstantially, etc.) possible worlds. Following Cariani et al., I'll call this set of worlds the "modal background" or—because my focus will be information-sensitive modals—the "information state". The ordering source is also a set of propositions, which determines a partial ordering over worlds in terms of some sort of ideality: for example, moral ideality, plausibility, etc. These two parameters determine the domain of the modal.

I'll gloss the classic account in terms of deontic modals for convenience. (The account is uniform across different flavors of modality.) Let a modal background i be a set of (for our purposes, epistemically) possible worlds and the ordering source d determine a deontic partial ordering in terms of some sort of ideality over those worlds.[1]

Definition 1. $w \leq_d w'$ iff according to d, w is at least as ideal as w'.

> In Kratzer's semantics, w is at least as ideal as w' iff $w \leq_d w'$ iff $\{p \in d : w' \in p\} \subseteq \{p \in d : w \in p\}$

Definition 2. $O_{w,i,d} =_{\text{df}}$ the set of worlds $w \in i$ s.t. for all $w' \in i$, if $w' \leq_d w$ then $w \leq_d w'$.[2] $O_{w,i,d}$ is the domain of the modal.

Kratzer's account of conditionals is built from pairing the classic account of modals with the restrictor analysis of conditionals. Where '$\Box$' is read *ought, should, must,* etc.,[3]

Modals: $\ulcorner \Box\phi \urcorner$ is true at a triple $\langle w,i,d\rangle$ iff ϕ is true at all worlds in $O_{w,i,d}$.

Conditionals: $\ulcorner$if $\phi, \Box\psi\urcorner$ is true at $\langle w,i,d\rangle$ iff $\ulcorner \Box\psi \urcorner$ is true at $\langle w,[i+\phi],d\rangle$, where $[i+\phi] = i \cap [\![\phi]\!]^{w,i,d}$.

2.2. A counterexample

A recent literature on deontic modals (Kolodny & MacFarlane 2010, followed by Cariani, Kaufmann, & Kaufmann 2011 and Charlow 2012) has focused on a class of counterexamples to Kratzer's semantics. Here is a simple counterexample:

[1] This is a simplification: in Kratzer's account, i and d are functions from worlds to modal bases and ordering sources. Throughout this paper I'll ignore this complication for ease of exposition.

[2] I presuppose the limit assumption for convenience.

[3] Throughout the paper I will ignore complications involving differences between strong and weak necessity modals.

The Three Envelope Puzzle

> You may have your choice of one of three envelopes: A, B, and C. Either envelope A or envelope B contains $10, but you have no idea which. The other is empty. Envelope C contains $9. The contextually salient priority is to get as much money as you can.[4]

The following sentences are judged true in this context.

(1) a. You should take envelope C.
b. If the $10 is in envelope A, you should take envelope A.
c. If the $10 is in envelope B, you should take envelope B.

Kratzer semantics predicts, falsely, that these sentences are inconsistent in this context.

The priority in this case is to gain as much money as possible. So suppose the ordering source simply orders worlds by how much money you receive:

(2) $d = \{$you gain $1, you gain $2, ..., you gain all the money in the world$\}$

And so on this ordering, the domain $O_{w,i,d}$ is the set of worlds in the modal background where you are lucky and choose whichever envelope the $10 happen to be in. But if this is the ordering, then (1a) cannot be true. So certainly this ordering doesn't get the right result. Indeed, this ordering entails that (3) is true:

(3) You should choose A or B.

You might be thinking: what went wrong is that the ordering we used didn't involve the sort of priority that is sensitive to ignorance: in this case, ignorance about the location of the $10. What we need is an ordering determined by an information-sensitive priority. So, perhaps a better ordering is one that orders worlds according to the *expected* monetary value. And this does allow us to get the right result for (1a): the worlds where you maximize monetary value are all worlds where you choose envelope C.

Be that as it may, Kratzer semantics only gets the right prediction for (1a) at the cost of getting the wrong prediction for (1b) and (1c). We are now forced to predict that both are false, relative to this ordering.

Consider (1b): the antecedent eliminates all of the worlds where the $10 isn't in envelope A. But because the ordering is fixed independently of the modal background, its ranking of worlds has to remain unchanged. Since there are still worlds in the restricted modal background where we could

[4] What I'm calling the "Three Envelope Puzzle" is a variation, via (Ross 2012), on the Miners Puzzle, discussed in Kolodny & MacFarlane (2010), who take this example from Parfit (unpublished).

take envelope C, those are predicted to be the best worlds in the restricted modal background. So with this ordering, Kratzer semantics incorrectly predicts (4a) and (4b) to be true.

(4) a. If the $10 is in envelope A, you should choose envelope C.
b. If the $10 is in envelope B, you should choose envelope C.

But, it might be protested, conditional on $10 being in envelope A, surely the action with the greatest expected monetary value is choosing A. And that's obviously true. The problem is that in order to represent this, we need the ordering of worlds to be shiftable with changes in the modal background. And Kratzer semantics doesn't allow that. Kratzer semantics lets embedding under conditionals shift the modal background, but not the ordering. And examples like the Three Envelope Puzzle show that embedding under conditionals can shift the ordering as well.[5]

The problem here is not specific to the suggested orderings we've just considered. It's a general problem: there isn't *any* ordering that can be plugged into Kratzer semantics such that (1a), (1b), and (1c) are consistent. That's the fundamental problem for the Kratzer framework. Those three sentences exhibit a property that Kolodny & MacFarlane (2010) call "serious information dependence". Kratzer semantics incorrectly rules out the possibility of serious information dependence.

SERIOUS INFORMATION DEPENDENCE: Given some body of priorities and set of options, acquiring more information can change which of the available worlds are best.

Formally: There is some world w' in both a modal background i and a strengthening of that modal background $[i+\phi]$ such that w' is in $O_{w,i,d}$ but not best in $O_{w,[i+\phi],d}$.[6]

Cariani, Kaufmann, & Kaufmann (2011) offer a tidy general proof that Kratzer semantics can't allow serious information dependence:

> Suppose for reductio SERIOUS INFORMATION DEPENDENCE. So there's some $w' \in [i+\phi], [i+\phi] \subseteq i$, such that (i) $w' \in O_{w,i,d}$ but (ii) $w' \notin O_{w,[i+\phi],d}$. Because of (ii) and Definition 1, there's some $w'' \in [i+\phi]$ such that $w'' \leq_d w'$ but $w' \not\leq_d w''$. But then $w'' \in i \supseteq [i+\phi]$. And so $w' \notin O_{w,i,d}$.

As a result, in order for Kratzer semantics to yield the correct prediction in this case, the context between (1a) and (1b) would have to differ. In other words, the sentences would have to exhibit some sort of equivocation. This is a bad result. Given that nothing new is learned and given that the

[5]To be clear, Kratzer semantics is perfectly compatible with the consistency of *some* sentences of the form: ⌜should ψ⌝ and ⌜if ϕ, should not ψ⌝. This can happen, for example, when the new modal base, restricted by ϕ, doesn't contain any ψ-worlds, as well as in other cases. The problem is when some of the very same worlds are still available in the restricted modal base, but go from being ideal to nonideal.

[6]Cariani, Kaufmann, & Kaufmann (2011) and Charlow (2012) give a more thorough explanation of why Kratzer cannot accommodate SERIOUS INFORMATION DEPENDENCE.

sentences themselves seem to contain no indicators that the relevant priorities differ, there's just no reason to expect that there should be a change in context. Modal parameters aren't supposed to be infinitely flexible formal tools. If we allow for appeals to shifts in the parameter where there's no evidence of shifts in priorities, our metasemantics is going to be too flexible to be predictive.

There is, by contrast, strong evidence that deontic modals shouldn't be expected to behave as the Kratzer account suggests. Here is a simple argument for the claim that deontic modals are seriously information dependent.

The simple argument

Our semantics for deontic modals shouldn't preclude the possibility of expressing, without equivocation, the consequences of a reasonable decision rule. So, for example, it shouldn't rule out the possibility that, in some contexts, the salient priority is to maximize the expectation of some kind of value (money, hedons, strawberries, etc.). There might in some contexts be a reading of *should* where, e.g., *John should* ϕ is true iff ϕ*ing maximizes expected x* is true. On this reading, any world (in the modal base) where John maximizes expected x is ideal within the modal base; in other words, a world in $O_{w,i,d}$. So $[\![\text{John maximizes expected } x]\!]^{w,i,d} \cap i = O_{w,i,d}$. The reading we're after makes *should*, like expected value, vary with a body of information. That body of information can be picked out by the modal background parameter.

All of this should be uncontroversial. Now, a piece of data: this reading of the deontic necessity modal allows for the consistency of *John should* ϕ and *If* ψ*, John should not* ϕ even though ϕing is still an option. It follows from the consistency of ϕ*ing maximizes expected value* and *if* ψ*, then not* ϕ*ing maximizes expected value* even though ϕing is still an option. And so this is enough to generate serious information dependence.

The suggestion here is not that an adequate semantic account should incorporate a disambiguation of *should* such that *should* ϕ is true iff ϕing maximizes expected utility. The suggestion is rather that the semantics for deontic modals shouldn't rule out the possibility that maximizing expected utility is the sole salient priority in some contexts. We shouldn't build decision theory into our semantics, but we also shouldn't make the semantics incompatible with even *expressing the consequences* of a decision theory. And that's what Kratzer semantics effectively does.

2.3. The ellipsis hypothesis

Let me briefly address an attempt to avoid the problem, discussed by von Fintel (2012), who cites unpublished notes by Kratzer.[7] An idea that is commonly floated is that we can explain away the pattern of behavior I've described for deontic modals and other information-sensitive expressions

[7]Others have independently suggested this hypothesis in personal communication.

by denying the literal and non-elliptical truth of (1b) and (1c). On this view, those two conditionals must elliptically express (5a) and (5b) in order to be true:

(5) a. If the $10 is in envelope A *and you know/learn that it is*, then you should choose envelope A.
b. If the $10 is in envelope B *and you know/learn that it is*, then you should choose envelope B.

This is a big shift from the familiar accounts. It's broadly accepted that ⌜*if* ϕ, ψ⌝ is not equivalent to ⌜*if* ϕ *and we learn/know* ϕ, ψ⌝.

I will briefly note two problems with this account. One problem is that in some Three Envelope Puzzle contexts, (5a) and (5b) suffer presupposition failure. The Three Envelope Puzzle sentences are still all true when we stipulate that the case is one where we *know* we don't know and won't learn where the $10 is before the choice must be made. But in the context where it's ruled out that we'll learn the $10's location, the antecedents of (5a) and (5b) are incompatible with the speakers' knowledge and beliefs and with the context set (the set of worlds treated as possible according to conversational participants' common knowledge). Indicatives presuppose that their antecedents are compatible with the context set (see e.g. Stalnaker 1975) or at least the speaker's knowledge or beliefs. So (5a) and (5b) have a false presupposition. But the Three Envelope Puzzle conditionals (1b) and (1c) show no signs of suffering presupposition failure.

Second, note that conditionals don't generally allow for this sort of addition of *and we learn/know it* to the antecedent. Compare:

(6) a. If my partner is planning a surprise party for me, I should try not to find out that she is.
b. #If my partner is planning a surprise party for me and I learn/know that she is, I should try not to find out that she is.

And in fact, we can generate cases analogous to the Three Envelope Puzzle where it's clear that the conditionals ⌜*if* ϕ, ψ⌝ cannot be glossed as ⌜*if* ϕ *and we learn* ϕ, ψ⌝.

> Suppose I prefer not to know what my partner got me for my birthday. If an action will probably ($> 50\%$ likely) lead to me knowing what she got me, I don't want to do it. (Of course, if I already know what she got me, this consideration will be moot.) I also want to get my scarf, which is in the closet. The present would be in the closet if and only if she got me clothing; but that's very unlikely.

Take a desire-based, instrumental reading of *should*, where the only salient priorities are:

(7) {I don't do anything that would probably lead to me learning what present my partner got

me; if it won't lead to me probably learning what present my partner bought got, I get my scarf}

The following sentences are true:

(8) a. I should look in the closet.
b. If my partner got me clothing, I shouldn't look in the closet.

But the Kratzer picture can't predict their compatibility in this context.

	clothes ($Pr < .5$)	not clothes
get scarf	w_1	w_2
don't get scarf	w_3	w_4

What's going on seems to be this: the case is one in which w_1 through w_4 form the modal background. w_1 and w_2 are better than w_3 and w_4, because the former satisfy both ordering source (OS) propositions—where the *probably* is assessed relative to the modal background and stipulated probabilities—and the latter only satisfy one of the OS propositions. But when the *if*-clause in (8b) restricts the modal base to w_1 and w_3, the *probably* is assessed relative to the restricted modal base and probability function conditionalized on the *if*-clause. There, w_1 satisfies only one of the OS propositions and w_3 satisfies both (the latter trivially). So the ordering of w_1 and w_3 is reversed—which, of course, the Kratzer semantics rules out.

How does the elliptical-*know/learn* hypothesis fare with this example?

(9) a. I should look in the closet.
b. If my partner got me clothing *and I know/learn it*, I shouldn't look in the closet.

Clearly (9b) is not equivalent to (8b); in this case, the former is false. So the elliptical-*know/learn* hypothesis cannot explain away these sorts of cases.

In what follows, I'll briefly discuss a proposed conservative amendment to the Kratzer framework (section 3). The problems that this proposal faces are useful to put on the table, because they help to expose what I think is a more fundamental problem with the Kratzer account and its variants (section 4). Understanding how these problems work will help to clarify the motivations for my own proposal (section 5).

3. A conservative solution?

Cariani, Kaufmann, & Kaufmann (2011) and Charlow (2012) independently give conservative amendments to the classic account of modals that involve similar basic operations: they introduce a third parameter that under certain circumstances coarsens the ordering source. In the Kratzer semantics, the ordering source ordered worlds, not options. These accounts order options. Doing so allows them to predict the consistency and truth of the Three Envelope Puzzle sentences. In what follows, I will focus on the example of the Cariani, Kaufmann, & Kaufmann account.[8]

Cariani, Kaufmann, & Kaufmann's third parameter is a "decision problem": roughly, a partition over the modal background such that each cell of the partition represents a different option that is choosable for the agent. An action α is **choosable** iff there's some action specification β such that it's epistemically necessary that the agent can knowingly perform β and it's epistemically necessary that performing β entails that α is achieved. So, for example, in the Three Envelope Puzzle scenario, the option of receiving \$9 is choosable. The option of receiving \$10 is not choosable: there's no action such that you can knowingly perform that action and that you know will lead to your receiving \$10.

As with Kratzer, the ordering is projected from a set of propositions. But whereas Kratzer defined the ordering in a fine-grained way, over worlds, Cariani, Kaufmann, & Kaufmann coarsen the ordering so that worlds are only ranked differently if they are in different cells of the partition of options. Kratzer's orderings, recall, are determined as follows:

(10) $w \leq_d w'$ iff $\{p \in d : w' \in p\} \subseteq \{p \in d : w \in p\}$

while Cariani, Kaufmann, & Kaufmann's orderings are projected from both the ordering source and the decision problem (where $[w]_\pi$ is the cell (option) of a partition π that w is in):

(11) $w \leq_d^\pi w'$ iff $\{p \in d : [w']_\pi \subseteq p\} \subseteq \{p \in d : [w]_\pi \subseteq p\}$

In other words, even if an option can have many possible outcomes with different objective values, the ordering of worlds doesn't reflect those differences: it only orders *options*, not outcomes.

In the Three Envelope Puzzle case, for example, suppose the ordering source is {you receive \$1, you receive \$2, ...}. Now, what you should do is be in a cell that entails as many of these propositions as possible. But since the cells are determined by which actions are choosable, there's

[8] Charlow's account is susceptible to many, though not all, of the objections I will lodge against Cariani, Kaufmann, & Kaufmann. One interesting feature of Charlow's account is that he believes that sentences like those in the Three Envelope Puzzle are inconsistent when *should* is replaced by the stronger deontic modal, *must*. The account he eventually gives has surprising implications for that distinction: in particular, his account seems to make strong deontic necessity modals neither stronger nor weaker than weak deontic necessity modals. I don't share his judgment (see also von Fintel 2012), and I find the result I mentioned unattractive. My eventual account has nothing helpful to say about the comparative felicity of strong versus weak necessity modals.

no cell that represents the action of receiving \$10. The partition, instead, will be: {you choose envelope A, you choose envelope B, you choose envelope C}. The cells where you choose either envelope A or envelope B include worlds where you receive \$10 and worlds where you receive \$0. Since they include the latter, they entail none of the ordering source propositions. But the cell where you choose envelope C entails that you receive \$9, and so entails nine of the ordering source propositions. So the best option according to our ordering is to choose envelope C, and so that option is the domain of the modal. So the Cariani, Kaufmann, & Kaufmann account makes the correct prediction for (1a).

When the modal base is restricted to worlds where the \$10 is in envelope A, the cell where you choose A includes only worlds where you receive \$10, and so entails ten ordering source propositions. Choosing C only entails nine (a proper subset of those entailed by choosing A). So Cariani, Kaufmann, & Kaufmann also make the correct prediction for (1b).

4. The deeper problem

Kratzer, in unpublished notes quoted by von Fintel (2012), argues that Cariani, Kaufmann, & Kaufmann's account, however conservative, involves some unnecessary addition of decision theoretic machinery: in particular, the decision problem parameter. Kratzer asks rhetorically: "Why pack information about rational decision making into the meaning of modals?" (von Fintel 2012, p. 25)—the implicature being, of course, that we shouldn't.

I agree with Kratzer that we shouldn't pack information about decision theory into the meaning of modals. But the Kratzer account, the Cariani, Kaufmann, & Kaufmann account, and the Charlow account *all* do. And they don't pack in just any information: rather, they include substantive normative assumptions, in the form of controversial decision theoretic rules. In section 4.1, I'll develop some problem cases for all three accounts and then in section 4.2, I'll explain how these cases are manifestations of normative commitments that are built into Kratzerian assumptions about modals and conditionals. For that reason, conservative amendments to the Kratzer picture will not yield an adequate account of information-sensitive deontic modals.

4.1. Problems with accounting for probability and value differences

Consider a variation on the Three Envelope Puzzle where instead of having no idea where the \$10 is, you're 97% confident that it's in A. Nothing else is changed: priorities are held fixed, and the same options are available; only your information has changed. If you choose envelope A, there's a very good chance you'll receive \$10. (12) is true:

(12) You should choose envelope A.

Cariani, Kaufmann, & Kaufmann are unable to predict (12) (relative to the fixed set of priorities). The problem: their account coarsens the ordering over worlds so that it only distinguishes between

chooseable outcomes: that is, outcomes that you know you can bring about. So they predict deontic modals to be insensitive to the probabilities of outcomes given our actions.

On Cariani, Kaufmann, & Kaufmann's account, any world w where you choose A is in the same cell of the decision problem partition as a world w' where you choose A and \$10 is in B. After all, that's still an epistemic possibility for you. w' is a world where you receive \$0. So the option of choosing A does not entail any of the propositions in the ordering source. But the option of choosing C still entails receiving \$9 (i.e. entails nine ordering source propositions). And so by their ranking, any world where you choose C is better ranked than any world where you choose A; so (1a) is still predicted to be true.

Kratzer's account exhibits the same insensitivity to adjustments of probability. In our example, the modal background still contains worlds where the \$10 is in A and worlds where it's in B. (The latter are just less probable.) And so worlds where it's in A and you choose A are no better than worlds where it's in B and you choose B. So (12) is predicted to be false: some worlds in the domain are choose-B worlds. It doesn't matter how improbable they are.

The coarsening account, like Kratzer semantics, is also not sensitive to cardinal differences in the value (desirability, moral status, etc.) of outcomes. These accounts all only consider the ordering, not whether one outcome is *a lot* better than another. Consider a case just like the original Three Envelope Puzzle except envelope C contains \$1 instead of \$9. The coarsening account predicts that (1a) will still be true:

(1a) You should choose envelope C.

While the option of choosing A or B doesn't entail any of the ordering source propositions, choosing C entails one of them. So Cariani, Kaufmann, & Kaufmann predict (1a) to be true. And if Kratzer were to predict the truth of (1a) in the original Three Envelope Puzzle case, relative to these priorities, she would have to predict it here too, at least on any natural ordering source.

But (1a) is false. In that circumstance, the best thing to do is to choose A or B at random.[9]

Of course, all three accounts could easily devise ad hoc ordering sources that happened to give the correct predictions for individual cases. But for any alternative ordering source that they suggest, we can generate counterexamples that have the same structures. The reason is that cardinal differences in probability and value matter for many of our normative judgments. So we shouldn't rule out the possibility of their mattering in the semantics of deontic modals.

[9] If this judgment is unclear, change the scenario so that there is \$1000 in A or B.

4.2. Normative commitments built into the semantics

These empirical objections to the Kratzer and Cariani, Kaufmann, & Kaufmann accounts point to a more general objection. Both accounts incorporate substantive normative assumptions into their semantics. And because these normative assumptions aren't very attractive and aren't the sorts of assumptions that guide ordinary speakers' judgments, they lead to a wide variety of incorrect predictions.

Cariani, Kaufmann, & Kaufmann rule out the truth of ⌜should ϕ⌝ if and only if the ϕ-worlds in the modal background include even one world with an objectively worse outcome than the worst outcome possible for some alternative to ϕ. Normatively, this amounts to the assumption that an action α is worse than an action β if and only if the worst possible outcome of α is worse than the worst possible outcome of β.

In other words, Cariani, Kaufmann, & Kaufmann encode the decision rule Maximin into the semantics of deontic modals.[10] This is not a good thing to build into our semantics.[11]

So I agree with Kratzer that their account builds decision theoretic information into the meanings of normative language. (Though she seems to think that merely the partition of options is problematic, and I don't; I take issue instead with the normative commitments.) But in fact, Kratzer's account is guilty of the same charge. Whereas the Cariani, Kaufmann, & Kaufmann account encodes the decision rule Maximin, Kratzer's account encodes Maximax: roughly, the rule that one should choose the option that has some chance of having the best possible outcome. This is a straightforward consequence of the more basic commitment of that semantics: that we should always simply bring about the best possible outcome in the modal background.[12]

Note: I am not suggesting that the relevant commitments are commitments to norms of maximizing the maximum or minimum quantity of *utility*. Maximax and Maximin are not strictly norms relating to utility. They are better understood as rule schemas, requiring that we maximize the maximum or minimum quantity of *some sort of value*, which here is determined by the contextually selected ordering source. It might be money, or strawberries, or quantities of money >\$7, etc. Variations in the sort of value at issue will vary the predictions of the accounts. But any time the source of value imposes a total ordering over worlds in the modal background, these commitments will manifest themselves for each account.

The commitments to Maximax and Maximin in the two accounts I discussed can't be written off

[10] Charlow's semantics, which I earlier mentioned faced many of the same objections as the CKK account, avoids this charge. (His normative view entails Maximin, but it isn't semantically encoded.) But his semantics *does* rule out other decision rules—in particular, decision rules that are sensitive to probabilistic information and cardinal value differences.

[11] Thanks to Paolo Santorio for discussion on this point.

[12] Thanks to Fabrizio Cariani for discussion on this point.

as a byproduct of some idealization. Neither of the views under discussion allows deontic modals to be sensitive to probabilities. And while decision rules like Maximax and Maximin show no sensitivity to probabilities, other decision rules do. The decision rules that best characterize how in ordinary circumstances agents make decisions—rules like expected value maximization—all require this form of sensitivity.

The account I propose will retain and extend the idea of letting context determine the priorities and information that determine the meanings of modals. Modals and conditionals are relativized to an information state of some sort and to some body of priorities. But instead of determining how these two interact by building decision theoretic norms into the semantic machinery, I suggest we let the relevant decision theoretic norms also be determined by context.

5. The proposal

5.1. Broad brushstrokes

The view I want to sketch will retain variations on the classic account's modal base and ordering source. We will need a parameter that plays roughly the role of modal backgrounds and another that plays roughly the role of ordering sources. In addition, though, we need a third parameter that is sensitive to uncertainty in an information state and to information about a body of priorities. The value function parameter, v, picks out contextually salient priorities; it is a variation on Kratzer's ordering sources. Our new parameter, by contrast, is sensitive to norms of rational decision-making under conditions of uncertainty. Call this third parameter r. While v looks for sources of value, r looks for rational decision rules.

What kinds of decision rules does r look for? There are plenty of alternatives: expected value maximization, Maximin, Maximax, some modification of expected value maximization that allows for rational risk aversion, rational sunk cost reasoning, or whatever. In all cases, the kind of value in question is provided by the priorities parameter, analogous to the ordering source.

This means we have two normative parameters. Why? There cannot be an ordering source that characterizes any of the most plausible candidates for a decision rule under uncertainty. There cannot be an ordering source that says, for example, maximize expected monetary value. An action only maximizes expected value relative to an information state. Ordering sources are, definitionally, not relativized to information states. If we want deontic modals to be able to reflect a ranking of worlds based on whether agents do something like maximizing expected value, and still preserve the observed facts about conditionals, then we need a ranking that is not independent of the information state.

Then why have an objective value parameter? Why not just some sort of function from information states to sets of worlds where the subjectively best actions are performed? This is the route that Kolodny & MacFarlane (2010) take. But they do nothing to explain what, in a particular case,

would determine the relevant function. So the account doesn't enable us to make predictions about which deontically modalized sentences are true; it only allows us to predict their consistency and inconsistency. (See (Charlow 2012) and (Cariani, Kaufmann, & Kaufmann 2011) for a fuller development of this objection.) I think the most plausible way of spelling out what determines a deontic selection function in a context will be a story of the form I'm characterizing: that is, a story according to which, in a particular context, the domain of quantification for deontic modals is determined by a function from what's valued, what's uncertain, and some decision rule to what's subjectively best.

5.2. The machinery

Apart from the addition of the third parameter, my account makes two substantial departures from Kratzer semantics:

I argued above that subjective deontic modals need to show sensitivity to probabilities of some sort. So just as we need to access some sort of probability function for the information-sensitive operator *probably*, the same will be true for deontic modals. I also argued that we should make our priorities-sensitive parameter potentially sensitive not just to ordinal ranking of different outcomes, but also to cardinal differences between outcomes. These two changes mean that we can no longer represent either of those two parameters as mere sets of propositions.

So, our three parameters:

1. An informational parameter s. Following Yalcin's (2010) suggestion for the semantics of *probably*, s will pick out pairs $\langle i, Pr\rangle$ of a modal background (a set of worlds) and a probability function.[13,14]

2. A value parameter v. This parameter will pick out some sort of value function, which we can treat as a function from worlds to real numbers.[15]

[13]Why isn't it redundant to use both a modal background and a probability function, rather than just considering the set of all worlds assigned positive probability by the probability function? First, I think we need to be able to shift the modal background to include and exclude information independently of the probability function. Second, as Yalcin notes, we want to allow the modal background to include possibilities that have probability zero (e.g. the possibility that if I throw a point-sized dart onto the dartboard, it'll hit exactly the point-sized center).

[14]Rothschild (2012) argues that in the case of *probably*, we actually need sets of probability functions in order to handle a puzzle from Swanson (forthcoming) about disjunctions of *probably* sentences. The puzzle is that there are true disjunctions of the form ⌜probably ϕ and probably ψ⌝ where neither disjunct is true. I am skeptical that this adequately motivates moving from probability functions to sets of probability functions. It seems to me that these are cases where the disjunction's *probably*s reflect some form of objective probability where one or the other disjunct is true (but the speaker doesn't know which). Neither disjunct is true only relative to the speaker's subjective probabilities or credences.

[15]This is an oversimplification. If we just use a unique value function, we lose one advantage that Kratzerian ordering sources have: the ability to represent incommensurability. I stick with a unique value function for the sake of simplicity of exposition, but a fully fleshed out version of this account will need *sets* of value functions in order to

3. A decision rule parameter r. r will pick out a set of decision rules.

The rational decision parameter, r, is similar structurally to a Kratzerian ordering source. Instead of sets of propositions, we can use sets of functions from $\langle s, v\rangle$ pairs to propositions. Call a set of such functions an r-ordering source. This amounts to an interpretation of the kinds of things that decision-theoretic norms are: they are norms that don't merely assess worlds, but instead assess worlds relative to bodies of information and priorities. For that reason, decision theoretic norms cannot be represented by ordering sources that are independent of the modal background.

We saw that ordering sources containing sets of propositions like those in (13) couldn't yield the appropriate kind of information-sensitivity:

(13) a. {You earn \$1, you earn \$2, ...}
b. {You maximize expected value}

But r-ordering sources—sets of functions from $\langle s, v\rangle$ pairs to propositions—can. For example, we can have ordering sources like these:

(14) a. {You perform an action that maximizes s-expected v-value}
b. {You perform an action that maximizes the minimum s-possible v-value}
c. {You perform an action that maximizes the maximum s-possible v-value}[16]

There can also be r-ordering sources where the $\langle s, v\rangle$ pair is idle in determining the relevant propositions.

(15) {You don't violate any federal laws $\langle s, v\rangle$} $\leftarrow$ $\langle s, v\rangle$ doesn't do anything.

Now, this new kind of ordering source allows us to define an ordering over worlds in the modal background, in almost the same way that Kratzerian ordering sources did.

Definition 3: $w \leq_{r,s,v} w'$ iff $\{p \in r_{s,v} : w \in p\} \supseteq \{p \in r_{s,v} : w' \in p\}$

> In English: w is at least as good as w' iff the set of $r_{s,v}$-propositions that w satisfies includes the set of $r_{s,v}$-propositions that w' satisfies.

Definition 4: The domain $O_{w,r,s,v}$ of a modal = $\{w \in i : \forall w' \in i, w' \leq_{r,s,v} w \rightarrow w \leq_{r,s,v} w'\}$

fulfill two roles: (i) representing incommensurability and (ii) allowing for sensitivity to cardinal differences in value.

[16]Why do all of these r-ordering sources only contain one decision rule? Only because no sets of decision theoretic norms that could all be contextually salient spring to mind. Note that if some decision theoretic norms are, in a context, prioritized over other decision theoretic norms, we might need sequences of r-ordering sources. This move is already at play in some versions of Kratzer semantics, e.g. (Rubinstein 2012).

> In English: the domain of the modal is the set of worlds in the modal background such that no world in the modal background is better than them (according to the ordering imposed by $r_{s,v}$).

Now, we update the classic semantics with our new parameter and the modified old parameters:

Modals: $\ulcorner\Box\phi\urcorner$ is true at $\langle w,s,v,r\rangle$ iff ϕ is true at all worlds the modal's domain, $O_{w,r,s,v}$

Conditionals: $\ulcorner$if $\phi,\Box\psi\urcorner$ is true at $\langle w,s,v,r\rangle$ iff $\ulcorner\Box\psi\urcorner$ is true at $\langle w,[s+\phi],v,r\rangle$,

> where $[s+\phi]=\langle i\cap[\![\phi]\!],Pr_\phi\rangle$, $Pr_\phi=Pr(\cdot\mid[\![\phi]\!])$[17]

It's easy enough to see how we can predict the joint truth of the Three Envelope Puzzle sentences when: (i) $r=\{$you maximize s-expected v-value$\}$; (ii) v assigns value to worlds according to money received; and (iii) $Pr(\$10\text{ in A})=Pr(\$10\text{ in B})=.5$.

There are many ways that the decision rule parameter might be resolved. First, it can be given explicitly: for example, *If Maximax is right, then we should choose A or B*. Second, the contextually salient r in some contexts might be whatever the "one true decision rule" is. Conversational participants might tacitly presuppose an objectively correct decision rule, in the same way they might presuppose that there is an objectively correct body of moral norms—even when these rules or norms aren't transparent to speakers. Finally, in a given context, which decision rule is relevant may be underdetermined. So which proposition is expressed by the sentences might also be underdetermined, and might get determinate truth values only supervaluationally. This seems to me a welcome result.

6. Conclusion

The account I've given is a generalization of Kratzer's semantics for modals and conditionals, like Kolodny & MacFarlane's account and Cariani, Kaufmann, & Kaufmann's account. The Cariani, Kaufmann, & Kaufmann account is still too strong: it constrains the choice of r to Maximin. The Kolodny & MacFarlane account, meanwhile, is too weak: it fails to be predictive. My account aims to find a happy middle.

[17] The amendments to the classic semantics that I have argued for—a third parameter for decision rules and an enrichment of the information and priorities parameters—are neutral about the question of whether modals are assessment-sensitive.

This new semantics abstracts away from complications, noted by Yalcin (2010) and Kolodny & MacFarlane (2010), about the behavior of conditionals with information-sensitive modals in their antecedents. I'm happy to take on board their suggested solution.

Note that, by design, this leads to the same sort of semantics for *probably* as Yalcin's account in (Yalcin 2010). All we need is an additional claim of the form: $\ulcorner$*probably* $\phi\urcorner$ is true at $\langle w,s,v,r\rangle$ iff $Pr(\phi)>.5$, or some more sophisticated modification.

I've shown that information-sensitive deontic modals exhibit serious information dependence. I've argued that one can give an account of deontic modals that predicts and explains their behavior without any substantive normative commitments. Other accounts that have been offered either fail to be predictive or incorporate unwarranted normative assumptions into the meanings of modals. The account I provide respects the motivations for the classic account of modals and conditionals, but excises the normative commitments latent in the classic account.

References

Cariani, Fabrizio, Stefan Kaufmann, and Magdalena Kaufmann (2011). "Deliberative Modality under Epistemic Uncertainty."

Charlow, Nate (2012). "What We Know and What to Do." *Synthese*, pp. –.

Kolodny, Niko, and John MacFarlane (2010). "Ifs and Oughts." *The Journal of Philosophy*, 107(3).

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.

Parfit, Derek (unpublished). "What We Together Do."

Ross, Jacob (2012). "Rationality, Normativity, and Commitment." *Oxford Studies in Metaethics*, 7: pp. 138–181.

Rothschild, Daniel (2012). "Expressing Credences." *Proceedings of the Aristotelian Society*, 112(1pt1): pp. 99–114.

Rubinstein, Aynat (2012). *Roots of Modality*. Ph.D. thesis, University of Massachusetts Amherst.

Stalnaker, Robert (1975). "Indicative Conditionals." *Philosophia*, 5(3): pp. 269–286.

Swanson, Eric (forthcoming). "The Application of Constraint Semantics to the Language of Subjective Uncertainty." *Journal of Philosophical Logic*.

von Fintel, Kai (2012). "The Best We Can (Expect to) Get? Challenges to the Classic Semantics for Deontic Modals."

Yalcin, Seth (2010). "Probability Operators." *Philosophy Compass*, 5(11): pp. 916–37.