THE *IF P, OUGHT P* PROBLEM

BY

JENNIFER CARR

Abstract: Kratzer semantics for modals and conditionals generates the prediction that sentences of the form *if p, ought p* are trivially true. As Frank and Zvolenszky show, for certain flavors of modality, like deontic modality, this prediction is false. I explain some conservative solutions to the problem, and then argue that they are inadequate to account for puzzle cases involving self-frustrating *ought*s. These cases illustrate a general problem: there are two forms of information-sensitivity in deontic modals. Even generalizations of Kratzer semantics that predict these two roles for information, e.g. Kolodny and MacFarlane predict that they vary together. I propose a generalization of Kratzer semantics that allows the two information roles to vary independently of each other.

Introductory modal logic told us that modals quantify over possible worlds: *ought p* is true iff *p* is true at all worlds within in a specified domain. And the popular restrictor analysis of conditionals told us that conditionals are a kind of modal: *if p, q* is true iff *q* is true at a specified domain of *p*-worlds. When these two views are paired, as in the familiar Kratzer semantics,[1] they have an unintended consequence: they validate:

(1) If *p*, ought *p*.

But obviously not all instances of *if p, ought p* are trivially true:[2]

(2) If you beat up elderly people, you ought to beat up elderly people.

This point was first made about deontic conditionals under the restrictor analysis by Frank (1997); Zvolenszky (2002, 2006, 2007) provides detailed investigation. The point generalizes to other analyses of deontic conditionals, and was independently discovered by van Fraassen (1972), Spohn (1975), and Jackson (1985).

Pacific Philosophical Quarterly •• (2014) ••–•• DOI: 10.1111/papq.12048

JOBNAME: No Job Name PAGE: 2 SESS: 66 OUTPUT: Mon Oct 6 17:21:13 2014 SUM: 5335E046
/Xpp84/wiley_journal_P/PAPQ/papq_v0_i0/papq_12048

If we want to hold onto possible worlds semantics for modals and a restrictor analysis of conditionals, how can we avoid this result? I will argue that doing so is harder than expected. Section 1 introduces the 'classic' account of modals and restrictor analysis of conditionals, and explains how they validate *if p, ought p*. In Section 2, I describe some candidate explanations for how *if p, ought p* can fail to be true that are consistent with the restrictor analysis combined with possible worlds semantics for modals. According to the most successful of these theories, iffy *ought*s are systematically ambiguous between a single-modal interpretation and a double-modal interpretation. Section 3 discusses a test for disambiguating the single- and double-modal readings, and shows that this test gives counterexamples to the proposed account. Section 4 surveys strategies for maintaining a conservative semantics while accommodating tough cases. Finally, in Section 5, I provide a positive account. I show that two forms of information-sensitivity affect the interpretation of iffy *ought*s. Generalizing our account of modals to allow them to vary independently of each other makes it possible to model and predict problem cases.

1. The problem

1.1. A SIMPLIFIED KRATZER SEMANTICS

Why is the schema *if p, ought p* predicted to be valid? Because the conditional is true iff in all the best possible *p*-worlds, *p* is true.

To show this more carefully, let's briefly rehearse the Kratzer account of conditionals, which has long been the default theory in linguistic semantics. According to Kratzer semantics, conditionals are a kind of quantified sentence; usually involving a modal quantifier. Modals, in this view, are given a possible worlds interpretation: *ought p* is true iff *p* is true at all the best worlds that are possible, given the circumstances.[3],[4],[5]

Two elements of this analysis are contextually determined: what's *best* and what's *possible given the circumstances*. There are different ways of modeling these two contextual parameters. Kratzer models them with two sets of propositions (i.e. sets of sets of worlds): a modal base, which is a set of propositions characterizing the circumstances (determining the set of relevant possible worlds), and an ordering source, which determines a pre-order over worlds.[6] (The preorder can be used to represent the relation of comparative deontic ideality, teleological ideality, epistemic plausibility, and so on.)

For ease of exposition, I will model these with simpler and more idealized tools: a set of worlds and a (total) ordering. (See Figure 1.) The modal background, *f*, is the set of worlds that are possible given the circum-

JOBNAME: No Job Name PAGE: 3 SESS: 66 OUTPUT: Mon Oct 6 17:21:13 2014 SUM: 54B64968
/Xpp84/wiley_journal_P/PAPQ/papq_v0_i0/papq_12048

stances (for example, epistemically possible worlds). The ordering, *g*, ranks worlds in terms of ideality of some sort or other; we'll focus on deontic ideality. These jointly determine the domain of the modal:

> DOMAINS: *domain*(*w, f, g*) is the set of worlds in the modal background *f* ranked highest by the ordering *g*.

The official statement of this simplified Kratzer semantics:

> MODALS: *ought p* is true at $\langle w, f, g\rangle$ iff *p* is true at all $w' \in$ *domain*(*w, f, g*).

Conditionals are, on this view, modals that have a restriction on the modal background. The *if*-clause restricts that set of worlds: *if p, ought q* takes the set of worlds in the modal background *f* and eliminates all the worlds where *p* is false. *ought q* is evaluated relative to the remaining set of worlds. In Figure 2, I represent restrictions diagrammatically by graying out a portion of the modal background.

Official statement:

> CONDITIONALS: *if p, ought q* is true at $\langle w, f, g\rangle$ iff *q* is true at *domain*(*w, f* + *p, g*) where '*f* + *p*' is shorthand for the intersection of *f* and the set of worlds where *p* is true.

1.2. THE PUZZLING CONSEQUENCE

From here, it's easy to see how CONDITIONALS validates *if p, ought p*, relative to any world-modal background-ordering triple. *if p, ought p* is

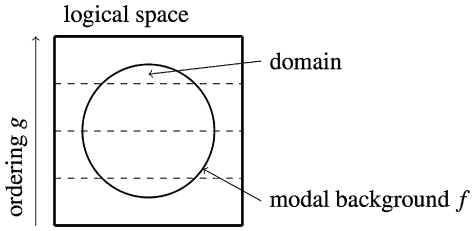

Figure 1 Modals.

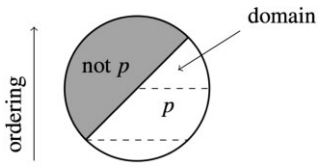

Figure 2 Conditionals.

JOBNAME: No Job Name PAGE: 4 SESS: 66 OUTPUT: Mon Oct 6 17:21:13 2014 SUM: 5FE62A4E
/Xpp84/wiley_journal_P/PAPQ/papq_v0_i0/papq_12048

true iff *p* is true in all of the *g*-ideal *p*-worlds that are possible according to *f*. It's trivial that *p* is true in all the *p*-worlds in any set of worlds.

When the necessity modal is epistemic, the triviality of *if p, ought p* is unproblematic:[7],[8]

(3) If it's raining, it must be raining. ✓
(4) If she's not home, she must not be home. ✓

Similarly, as Zvolenszky (2006) notes, with appropriate context, related teleological examples may be interpreted as trivial. They are uninformative but not false:

(5) a. A: What do I have to do to go to Berlin?
b. B: To go to Berlin, you have to go to Berlin. ✓
c. B: If you want to go to Berlin, you have to go to Berlin. ✓

But when the modal is interpreted as deontic, there are clear counterexamples:

(6) If you beat up elderly people, you ought to beat up elderly people. ✗
(7) If John spills wine, he should spill wine. ✗

Similarly for circumstantial modals.[9]

(8) a. If the coin lands heads, the coin has to land heads. ✗
b. (Equivalently) If the coin lands heads, the coin can't not land heads. ✗

Perhaps the clearest examples involve nomological necessity. Suppose the trajectory of a silver ion after passing through a Stern-Gerlach magnet is presupposed to be indeterministic. Then the following utterance, interpreted as a circumstantial modal, is false:

(9) If the silver ion veers upward, then it must veer upward. ✗

The *if p, ought p* problem is even more general. Sentences of the following forms are also made trivially true:

(10) If *p*, may *p*.[10]
(11) It's not the case that if *p*, not ought *p*.
(12) It's not the case that if *p*, may not *p*.

JOBNAME: No Job Name PAGE: 5 SESS: 66 OUTPUT: Mon Oct 6 17:21:13 2014 SUM: 5155A271
/Xpp84/wiley_journal_P/PAPQ/papq_v0_i0/papq_12048

These are equally unattractive results. (The reader may generate her own counterexamples.)

I'll focus on Kratzer's restrictor analysis because it is widely accepted in linguistic semantics, and because generalizing the discussion to other theories of conditionals would interfere with the exposition. But it's worth noting that these and related problems afflict many other theories of conditionals, including the material conditional analysis and the Stalnaker semantics. Explaining why will require more setup, so I put it off until Section 2.2. The focus, going forward, will be Kratzer semantics. My central question will be: if we want to hold onto a restrictor analysis of conditionals and a possible worlds account of modals, how can we avoid these results?

2. *Some hypotheses*

2.1. HYPOTHESIS #1: THE DIVERSITY CONDITION

What's wrong with sentences like: 'If you beat up elderly people, you ought to beat up elderly people'? A first suggestion: let's consider this in conjunction with A. N. Prior's (1958) Samaritan Paradox:

(13) a. John should help the assault victim.
b. *Doesn't entail:* Therefore, there should be an assault victim.

It would, of course, be a bad result if (13a) entailed (13b). But the familiar semantics for modals generates this entailment, as long as (13a) and (13b) are evaluated relative to the same modal background and ordering. If all the worlds in the domain are worlds where John helps the victim, then trivially all the worlds in the domain include a victim.

The Samaritan Paradox suggests that deontic modals have a diversity condition on their modal backgrounds:

> DIVERSITY: *ought p* presupposes that the modal background includes both p and $\neg p$ worlds.

If there is a diversity condition on modal backgrounds (under some circumstances), then (13b) can't be evaluated at the kind of modal background where (13a) is most naturally assessed, i.e. a modal background that presupposes the existence of an assault victim. Instead, it must include the possibility of there being no assault victim. But if the modal background allows that there might be no assault victim, then on natural assumptions about what is most ideal, at least some worlds where there is no assault victim are better than worlds where someone is assaulted, but

JOBNAME: No Job Name PAGE: 6 SESS: 66 OUTPUT: Mon Oct 6 17:21:13 2014 SUM: 4F2FA958
/Xpp84/wiley_journal_P/PAPQ/papq_v0_i0/papq_12048

then John helps them. And so, (13b) will be evaluated relative to a different modal background from (13a), and at one where (13b) is correctly predicted to be false.

It is worth mentioning that the diversity condition also has some independent, intuitive plausibility, at least in the deontic case. It's plausible that deontic *ought* implies both *can* and *can not*.

As Frank (1997) first noted, a diversity condition can helps with *if p, ought p*: when the *if*-clause removes from the modal background all the $\neg p$-worlds, the modal background at which *ought p* is evaluated can't respect the diversity condition. So *if p, ought p* suffers presupposition failure.

Problem solved? Unfortunately, no: there are serious objections to this account. First: consider our example: 'If you beat up elderly people, you ought to beat up elderly people.' This sentence doesn't seem to suffer presupposition failure. It isn't judged non-truth-evaluable. It is judged straightforwardly false.[11]

Second, a diversity condition may help with *if p, ought p*, but it doesn't help with the corresponding *may*-conditional. If anything, the problem is worse: a diversity condition would validate all instances of both *if p, may p* and *if p, may* $\neg p$. But both have counterexamples.

Finally, we don't want to predict that *if p, ought p* is always defective. As Frank (1997) and Zvolenszky (2002, 2007) argue, there are cases where such sentences are judged both true and informative:

(14) a. John knows the route well, so if John turns on Exit 49, then he should turn on Exit 49.
b. Ryan is a workaholic, so if Ryan is taking a break, he has to take a break.

The diversity condition would make it the case that the sentences in (14) cannot be true: they are immediately defective. So we can't explain away the problem with validating *if p, ought p* simply by imposing a diversity condition on the modal background. At the very least, it can't be the whole story; we need to accommodate the truth of some instances of *if p, ought p*.

2.2. HYPOTHESIS #2: DOUBLE MODALIZATION

Let's consider an alternative proposal. The restrictor analysis is committed to conditionals being a form of restricted modal. The analysis assumes that when an overt modal appears in the consequent of a conditional, its modal background is the one that's restricted by the *if*-clause. For bare conditionals, there's a covert epistemic *must*:[12]

JOBNAME: No Job Name PAGE: 7 SESS: 66 OUTPUT: Mon Oct 6 17:21:13 2014 SUM: 5BA53E5F
/Xpp84/wiley_journal_P/PAPQ/papq_v0_i0/papq_12048

(15) a. If John spilled his wine, he's drunk.
b. If John spilled his wine, he must be drunk.

In other words: in some cases, the modals that *if*-clauses restrict are pronounced; in other cases, unpronounced.

Suppose instead that the covert modal is there whether or not there is an overt modal in the consequent. When an overt deontic modal appears, it is embedded under the covert epistemic modal. So a sentence like 'If it's raining, you should take an umbrella' actually has the form: *if p,* $\checkmark_e\checkmark_d q$ (where ✓ is a necessity modal like *ought* or *must*, and subscripted *e* and *d* designate epistemic and deontic modality, respectively). In order to evaluate the conditional, we look at *p*-worlds in the epistemic modal background and evaluate each pointwise for the truth of $ought_d\, q$. (See Figure 3 for a diagrammatic representation.)

How this helps with our *if p, ought p* problem: the restriction imposed by the *if*-clause restricts the epistemic modal, but not the deontic modal embedded under it. The deontic modal's modal background is not restricted. And so the antecedent *p* isn't automatically true at all of the worlds in the deontic modal's modal background.

We look at all points in the epistemic modal background where you beat elderly people, and evaluate each pointwise for the truth of *You ought to beat up elderly people*. *You ought to beat up elderly people* is true at an epistemically possible point iff *You beat up elderly people* is true in the best worlds that are possible, given the circumstances (i.e. in the deontic modal's unrestricted modal background). Let's see how this result is achieved in the case of (6).

(6) a. If you beat up elderly people, you ought to beat up elderly people.
b. If you beat up elderly people, it must be that you ought to beat up elderly people.

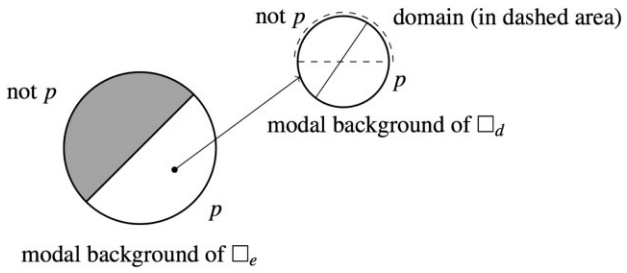

Figure 3 Double-modal conditionals.

But given the circumstances, it's possible for you not to beat up elderly people: even if you will in fact beat them up, you have the ability not to. So in the best worlds in the deontic modal background, you don't beat up elderly people. Therefore, it's not automatically the case that at each epistemically possible world where you beat up elderly people, you ought to beat up elderly people. So (2) is false.

It's easy to see how this also delivers the correct results for *if p, ought p* sentences that have true readings, by the very same reasoning. In these cases, the *if*-clause restricts the epistemic modal but not the deontic modal. So the deontic modal's modal background can still satisfy the diversity condition (if such a condition exists). And so this view can handle the objections to hypothesis #1.

There's a problem, though. If we model iffy *ought*s as including two modals, only one of which is restricted by the *if*-clause, we generate the wrong predictions for most ordinary instances of iffy *ought*s. A perfectly standard iffy *ought*:

(16) If you turn in your library books late, you have to pay a fine.

If we evaluate (16) as doubly modalized in the same way as (2), with plausible assumptions about the conversational background, it comes out false.

Take each epistemically possible world where you turn in your library books late, and evaluate whether 'you have to pay a fine' is true. The *if*-clause restricts only the epistemic modal and not the deontic modal. (This assumption is crucial for using the double-modal account to predict contingent instances of *if p, ought p*.) But because the deontic modal background can include *you-don't-turn-in-books-late*-worlds, it's not the case that you have to pay a library fine: in some worlds (indeed, the best worlds), you don't even have library fines. And so you don't have to pay library fines. The (16) is incorrectly predicted to be false. So the double-modal, single-restriction hypothesis isn't true, or at least isn't fully general.

The general question: how do we provide a unified account of conditionals like (16), where the antecedent acts as a restrictor on the deontic modal, and conditionals like (2), where the antecedent does not? This puzzle equally afflicts other analyses of conditionals, e.g. the material conditional account and the Stalnaker account. I will offer a very quick argument, for those familiar with the two views. (Other readers may prefer to skip ahead to the next subsection.)

First, the material conditional analysis doesn't predict any semantic interaction between *if*-clauses and deontic modals. So that view is ill-equipped to give an adequate semantics for run-of-the-mill iffy *ought*s like *If you're going to murder, you should murder gently*. The truth of that conditional is compatible with the truth of the negation of the consequent,

JOBNAME: No Job Name PAGE: 9 SESS: 66 OUTPUT: Mon Oct 6 17:21:13 2014 SUM: 534A265B
/Xpp84/wiley_journal_P/PAPQ/papq_v0_i0/papq_12048

You ought not murder gently, even at worlds where in fact the addressee is going to murder. The material conditional analysis can't be updated to incorporate the systematic restriction of embedded modals by the antecedent without thereby losing the analyses defining features, e.g. the validity of classical inference rules.

Second: Stalnaker semantics for conditionals is silent on the relation between modals and conditionals. But there are two basic strategies available to Stalnaker semantics: first, it could vary the selection function to a deontic selection function when there's a deontic modal in the consequent of the conditional. In that case, *if p, ought p* will be valid. So this is no progress. Second, it could retain a non-deontic selection function and simply evaluate deontic modal at the selected antecedent-world. In that case, Stalnaker semantics will make the wrong predictions for ordinary iffy *ought*s like (16). And so the Stalnaker account is in the same position as the Kratzer account.

2.3. HYPOTHESIS #3: HYBRID PROPOSAL

Hypothesis #2 suggested that iffy *ought*s have a covert epistemic modal with a deontic modal embedded under it. The *if*-clause restricts the epistemic modal rather than the deontic modal. But at least some iffy *ought*s require a single-modal reading, where the antecedent restricts the deontic modal. A natural hypothesis, then, is that both readings of iffy *ought*s are available. Conditionals of this kind are systematically structurally ambiguous, with both single-modal and double-modal readings.[13] As Geurts (2004) shows, an ambiguity of this sort is widespread in natural language.

While this looks to be promising, there's an immediate concern: on this view, 'If you smoke, you ought to smoke' would have to have a reading where it is trivially true. And it doesn't seem to. Similarly, 'If you turn in your library books late, you have to pay a fine' should have a false reading in the very same (natural) context where it has a true reading. If these secondary readings are available at all, they are certainly harder to access. It's worth asking whether there's a general explanatory model that would allow us to predict whether the single- or double-modal reading will be dominant for a given conditional.

Maybe this worry can be addressed. A naive hypothesis: double modalization is dispreferred and is typically accessible only when the singly modalized reading is defective.[14,15] According to this hypothesis, one circumstance in which the single-modal reading is defective is when it violates the diversity condition. And so sentences like 'If you smoke, you ought to smoke' are defective on the single-modal reading, but false on the double-modal reading. On the other hand, 'If Ryan is taking a break, he has to take a break' is defective on the single-modal reading, but *true* on the double-modal reading.

JOBNAME: No Job Name PAGE: 10 SESS: 66 OUTPUT: Mon Oct 6 17:21:13 2014 SUM: 53D30432
/Xpp84/wiley_journal_P/PAPQ/papq_v0_i0/papq_12048

And so we're able to predict that *if p, ought p* is neither trivially true nor trivially not true: some instances are true and some are false. So sentences of this form make substantial claims. (Meanwhile, more commonplace iffy *ought*s like (16) are not defective, and so are by default interpreted as only involving one modal.) In effect, we have taken the better of the previous two proposals. Each is individually inadequate, but together they make good predictions.

2.4. SINGLE-/DOUBLE-MODAL AMBIGUITY

Geurts (2004) shows that a similar ambiguity appears with other kinds of quantifiers. Conditionals are essential restricted quantifiers over possibilities (restricted modals). An example from Geurts' examples:

(17) If Beryl is in Paris, she often visits the Louvre.

Here the *if*-clause can restrict either the overt adverbial quantifier *often* or it can restrict a covert quantifier. Using the notation *Quantifier* [*restrictor*] [*nuclear scope*], we can spell out the two readings:

(18) a. Often [Beryl is in Paris] [Beryl visits the Louvre].
b. Must [Beryl is in Paris] [Often [] [Beryl visits the Louvre]]

On the single-quantifier reading, (18a), the *if*-clause restricts the overt quantifier *often*. (18a) says that among those occasions where Beryl is in Paris, many are such that she visits the Louvre. On the double-quantifier reading, (18b), the *if*-clause restricts a covert quantifier (plausibly epistemic *must*). (18b) says that among the epistemic possibilities where Beryl is in Paris, all are such that on many occasions she visits the Louvre.

If Geurts is right that this ambiguity is widespread among quantifiers, it's predictable that the single-/double-modal ambiguity would appear conditionals. The contrast between single-modal iffy deontic *oughts* and double-modal deontic iffy *oughts* seems to be a contrast between conditionals that express conditional obligations and conditionals that don't. In the latter cases, the antecedent merely provides evidence that supports the independent truth of the *ought* claim. Rough (!) glosses to exhibit the difference:

(19) If you spill wine, you have to mop.
a. *Conditional obligation gloss:* On the condition that you spill wine, you have to mop.
b. ≠ *Epistemic conditional gloss:* If you spill wine, then you must have to mop (already).

(20) If Ryan is taking a break, he has to take a break.
  a. *Epistemic conditional gloss:* If Ryan is taking a break, then Ryan must have to take a break.
  b. ≠ *Conditional obligation gloss:* Ryan has to take a break on the condition that he is taking a break.

A first-pass explanation of the distinction: in double-modal conditionals, the *if*-clause provides evidence for the obligation in the consequent. In single-modal conditionals, by contrast, the *if*-clause notes the circumstances where the obligation obtains. Single-modal deontic conditionals are typically used to express hypothetical imperatives or conditional obligations. Double-modal conditionals, by contrast, express something like speculation about possible obligations. For the moment, hopefully this rough characterization will suffice to distinguish the two types of iffy *ought*.

3. *Puzzle cases*

3.1. BACKGROUND ON INFORMATION-SENSITIVE DEONTIC MODALS

The problem that I will generate from the ambiguity account, and the solution that I propose, are tightly related to phenomena that central to the recent literature on the Miners Puzzle.[16] The kinds of information-sensitivity I will discuss share many of the features of the Miners Puzzle (though not those that have been the core of the recent debate on so-called 'serious information-sensitivity'). So it will be helpful to provide some background on the Miners' Puzzle phenomena.

The original Miners Puzzle, introduced to the philosophy of language literature in Kolodny and MacFarlane, 2010,[17] runs as follows:

> The Miners Puzzle.
> Ten miners are trapped either in shaft A or in shaft B, but we do not know which. [They are equally likely to be in either.] Flood waters threaten to flood the shafts. We have enough sandbags to block one shaft, but not both. If we block one shaft, all the water will go into the other shaft, killing any miners inside it. If we block neither shaft, both shafts will fill halfway with water, and just one miner, the lowest in the shaft, will be killed (Kolodny and MacFarlane, 2010, p. 115).

In this context, the sentences in (21) are all true:

(21) a. We ought to block neither shaft.
  b. If the miners are in shaft A, we ought to block shaft A.
  c. If the miners are in shaft B, we ought to block shaft B.
  d. The miners are either in shaft A or in shaft B.

Kolodny and MacFarlane argue that this case reveals two features of information-sensitive deontic modals, e.g. the so-called subjective *ought*. First, conditionals involving these modals can generate counterexamples to classical inference rules, e.g. (at least) proof by cases and modus tollens. (For arguments against alternative explanations for the apparent consistency of the sentences in (21), e.g. context-sensitivity or wide-scoping, see Kolodny and MacFarlane, 2010; Cariani, Kaufmann and Kaufmann, 2012; Carr, 2012; and Silk, 2014.) This fact about iffy *ought*s is consistent with Kratzer semantics.

Second, Kolodny and MacFarlane argue that this case reveals that *ought* is 'seriously information-sensitive': roughly, that the ordering of worlds at a context is a function of the modal background, rather than an independent parameter. In other words, changes in contextually salient information can alter not just which worlds are circumstantially possible, but which words are better than which. This fact about iffy *ought*s is not consistent with Kratzer semantics.

Why does the deontic ordering change with changes in the modal background? In (21a), *block-neither*-worlds are better than both *block-A*-worlds and *block-B*-worlds. When the modal background is restricted by an *if*-clause, as in (21b) and (21c), *block-neither*-worlds are worse than either *block-A*-worlds or *block-B*-worlds, even though all three types of world are still present in the restricted modal background. So the ordering must change.

An intuitive explanation for why iffy *ought*s generate both counterexamples to classical inference rules and serious information sensitivity:

The body of information that is relevant for information-sensitive modals is not always the speaker's, evaluator's, or conversational participants'. Under embeddings, the relevant body of information may often be a modified variant of the contextually salient information. In particular, in the consequent of a conditional, the salient information for an information-sensitive modal will be the speakers' information augmented with the information in the antecedent. For example, consider the following epistemically modalized conditional:

(22) If his lights are off, he must not be in his office.

This can be uttered in a context where no one has any idea whether he's in his office – i.e. where no one's body of information can be expressed by the consequent: *he must be in his office*. The modal instead quantifies over possibilities in the information state consisting of, say, the speaker's information augmented by the information in the antecedent.

The same phenomena affect other information-sensitive operators, e.g. *probably*:

(23) If his lights are off, he's probably not in his office.

And the same phenomena affect information-sensitive deontic modals. Again, in intuitive terms: in the Miners Puzzle (21a), there's a salient subjective reading of the *ought*. The salient priority determining the ordering is to aim to save as many miners' lives as possible, in light of the speaker's limited information. We might think of the salient priority as maximizing the expected number of miner's lives saved. (There are, of course, plenty of other norms that equally recommend hedging one's bets in cases where information is limited.) Suppose this is the very same priority in the Miners Puzzle conditional (21b). Suppose further that, as with (21), the salient body of information for the embedded modal is not the speaker's, but rather the speaker's information augmented with the information in the antecedent (that the miners are in shaft A). Relative to the augmented body of information, the subjectively best thing to do is to block shaft A. If the priority is to maximize expected miners' lives, then for the conditional, that expectation is relativized not to the speaker's actual credence function, but rather to their credence function augmented by (conditionalized on) the information in the antecedent. And so this is the option that maximizes the expected number of miners' lives. This is why the conditionals can be true relative to the same contextual parameters as the unconditional claim (21a).

Notice that this goes beyond the information-sensitivity suggested by Kratzer semantics. There, the antecedent affects the contextually salient information only by generating changes in the modal background. The Miners Puzzle shows that information from the antecedent can also affect the ordering of options, without any change in the contextual parameters. For example, if the prioritizing parameter is fixed by a contextually salient priority to maximize expected miners' lives, then the probability function used for the expectation is not the speaker's probability function, but rather the antecedent-augmented probability function. So the ordering generated by this priority shifts with linguistically generated information, not speakers' actual information.

So: the body of information relevant for information-sensitive operators embedded in the consequents of conditionals is not the contextually salient body of knowledge, but rather that information augmented by the information in the antecedent. This is true not just for modal backgrounds, but also for the deontic priorities that generate orderings for the modal. This conclusion is not entirely uncontroversial, but I won't defend it further.[18] I will assume its correctness for the puzzle cases I discuss below.

3.2. A NEGATIVE TEST FOR EPISTEMIC CONDITIONALS

In many circumstances it will not be clear whether the salient reading of an iffy *ought* is single- or double-modal. So it would be useful to have a test for single- and double-modal iffy *ought*s. Here is a (partial) candidate:

If iffy *ought*s are ambiguous, they generate ambiguity in argument forms. If you give a modus tollens-style argument with an iffy *ought*, your argument will be ambiguous between the following two forms.

Modus tollens
if p, $\Box_d q\neg$
$\Box_d q$
Doesn't entail: ¬ p
Quasi-modus tollens
if p, $\Box_e\Box_d q$
¬ $\Box_d q$
Does entail: ¬ p

On the standard semantics, epistemic conditionals, including double-modal iffy *ought*s, respect quasi-modus tollens (though not modus tollens).[19] As the literature on the Miners Puzzle has emphasized (especially Kolodny and MacFarlane, 2010), single-modal iffy *ought*s need not respect modus tollens.[20,21]

Assuming that iffy *ought*s can have a covert epistemic modal, the two arguments might be expressed in English in exactly the same way. So, a *prima facie* plausible hypothesis: we can learn that an iffy *ought* is not an epistemic conditional if it appears to violate modus tollens-style reasoning.

As the name suggests, quasi-modus tollens is *not* actually modus tollens. For this sort of conditional, the second premise would be the weaker claim: ¬ $\Box_e\Box_d q$. Epistemic conditionals need not respect modus tollens, but *prima facie*, they do seem to respect quasi-modus tollens.

Modus tollens
If it's raining, the streets must be wet.
The streets might not be wet.
Doesn't entail: Then it's not raining.
Quasi-modus tollens
If it's raining, the streets must be wet.
The streets are not wet.
Does entail: Then it's not raining.

Now, the most predictive account we have so far combines the single-/double-modal ambiguity with a diversity condition on deontic modal backgrounds. That account predicts that there are no true, single-modal instances of *if p, ought p*.

Can we find true instances of *if p, ought p* that generate violations of MT-style reasoning? If so, then according to our test, that account won't

JOBNAME: No Job Name PAGE: 15 SESS: 66 OUTPUT: Mon Oct 6 17:21:13 2014 SUM: 5179B0E2
/Xpp84/wiley_journal_P/PAPQ/papq_v0_i0/papq_12048

work. Such cases can't be double-modal because quasi-MT is (we hypothesize) valid. And they also can't be single-modal because they would violate the diversity condition.

3.3. STRUCTURAL AMBIGUITY + DIVERSITY + QUASI-MT: COUNTEREXAMPLE

We can find such cases in decision theory. I will focus on cases that have the same format as the Miners Puzzle, that is, on cases that generate violations of proof by cases. The information-sensitive modal behaves in the way that, I argued in Section 3.1, is standard for information-sensitive modals: in the consequent of the conditional, the modal is evaluated relative to the information of the consequent augmented by the information from the antecedent.

> Self-reinforcing *ought*
> Suppose your norm for behavior is to satisfy the desires you anticipate having in the future: roughly, to act in such a way that you think you'll be glad you did. You are in some decision situation with more than one option such that, predictably, if you choose that option, some rationalization process will kick in such that you'll retroactively desire that you chose that option.[22]

(24) a. If you choose the tapas restaurant, you ought to choose the tapas restaurant.
b. If you choose the Indian restaurant, you ought to choose the Indian restaurant.
c. You may choose the tapas restaurant and you may choose the Indian restaurant. (It's not the case that you ought to choose the tapas restaurant and it's not the case that you ought to choose the Indian restaurant.)
d. You will choose the tapas restaurant or the Indian restaurant.

Our test for epistemic conditionals suggests that (24a) and (24b) should be interpreted as conditional obligations and so it must be given a single-modal reading. After all, quasi-modus tollens fails. But the diversity condition predicts that on the single-modal reading, (24a) and (24b) suffers presupposition failure.

So, given that the sentences in (24) are true, they can't be interpreted as a single-modal iffy *ought*: this violates the diversity condition. But they also can't be interpreted as a double-modal iffy *ought*: this violates quasi-modus tollens.

3.4. STRUCTURAL AMBIGUITY + QUASI-MT – DIVERSITY: COUNTEREXAMPLE

One might say: perhaps we can get rid of the diversity constraint. After all, it's not so bad to allow that there might be recessive trivially true readings of sentences like *if you smoke, you should smoke*. Ordinary language users will tend to interpret contributions to the conversation as being informative (not trivially true), and so the contingent, double-modal reading will be salient.

But biting the bullet on this matter won't help: it can only explain self-reinforcing cases like (24), but not self-frustrating cases. We can generate other sorts of counterexamples to other schemas that the restrictor account, combined with possible worlds semantics for modals, makes false predictions about. For example, all instances of *if p, ought not p* are predicted to be trivially false.

A counterexample, again with a structure similar to the Miners Puzzle examples, and again taken from decision theory (via Gibbard and Harper, 1978):

> Self-frustrating *ought*
> If you are in the same city as Death tomorrow, then you'll die. Death has planned to be wherever he predicts you'll be, and he's very reliable in such predictions. Your options are to stay in Damascus or to go to Aleppo. But, as you know, if you stay in Damascus, then that's excellent evidence that Death will already be there. Similarly for going to Aleppo.[23]

Once again, we can generate a violation of modus-tollens-like reasoning:

(25) a. If you go to Aleppo, you ought not go to Aleppo (because Death will be there).
b. If you stay in Damascus, you ought not stay in Damascus (because Death will be there).
c. You may go to Aleppo and you may stay in Damascus. (It's not the case that you ought not go to Aleppo and it's not the case that you ought not stay in Damascus.)

Gibbard and Harper's discussion of this example in the decision-theoretic context provides a justification for accepting the conditional (25a): 'Any reason the doomed man has for thinking he will go to Aleppo is a reason for thinking he would live longer if he stayed in Damascus, and any reason he has for thinking he will stay in Damascus is a reason for thinking he would live longer if he stayed in Aleppo. Thinking he will do one is a reason for doing the other' (Gibbard and Harper, 1978, p. 156).

The conditional (25a) can't be a double-modal iffy *ought*: it would violate quasi-MT. And it can't be a single-modal iffy *ought*, for then the restrictor analysis predicts it to be trivially false.

The Death in Damascus example will be central for the remaining discussion, so it's helpful to understand its dialectical role. Some philosophers may want to challenge the data, typically by arguing that the conditionals (25a) and (25b) are false. I think these arguments often depend on false theoretic presuppositions rather than naked intuition. In particular, there is an entrenched but false theoretical presupposition that proof by cases is valid. But the Miners Puzzle shows that this is false; and independently, the standard semantics for conditionals never validated proof by cases.

The sentences in the Death in Damascus case are judged acceptable by many ordinary language speakers. But it's reasonable to question whether the data are the product of some kind of noise (for example, equivocation), or whether we should take them at face value. Instead of debating the data, let me offer what I take to be strong theoretical considerations in favor of the truth of the Death in Damascus sentences.

As I argued in Section 3.1, information-sensitive deontic modals pattern with other information-sensitive modals: the information relevant for the assessment of iffy *ought*s includes the information in the antecedent. The antecedent's information can have two effects: it can restrict a modal background, but it can also affect the deontic ordering. So, when the contextually salient priorities or norms are a function of information – for example, when they are expected utility theoretic – then embedded under conditionals, the relevant information for generating the ordering includes the information in the antecedent. Probability functions relevant for maximizing expectations are updated on the antecedent's information. This is the lesson of the Miners Puzzle.

When we apply this lesson in the Death in Damascus case, the sentences in (25) are exactly as expected. Suppose the contextually salient norm is a form of causal decision theory, where the Death in Damascus case was first introduced.[24] Then (25c) will be true iff going to Aleppo and staying in Damascus have equal causal expected utility. Suppose the *should*s in (25a) and (25b) behave in the same way as the Miners Puzzle conditionals do: the information from the antecedent is able to affect the ordering over worlds in the way that conditionalized probabilities affect the ordering of options in decision theory. That assumption, combined with the assumption that the relevant priority is to maximize causal expected utility, generates the prediction that *if p, ought q* will be true iff *q* uniquely maximizes causal expected utility relative to the relevant probability function conditionalized on *p*. So we expect (25a) and (25b) to be judged true at such a context, for the same reason that we should be able to predict the truth of (26):

(26) If you go to Aleppo, going to Aleppo does not maximize causal expected utility.

Now we just need a semantic framework that can accommodate this.

Finally, note that the aim here is not to give a decision theoretic semantics. I do not assume the truth of (any form of) causal decision theory, and I certainly don't assume that deontically modalized sentences express claims about causal decision theory. What I do assume is that natural language semantics for modals and conditionals is flexible enough to accommodate the various forms of causal decision theory and other normative systems that entail the truth of the Death in Damascus sentences without equivocation. If you don't accept the sentences, it should be for normative reasons. Their conjunction should not be ruled out as a matter of linguistic competence.

4. *Predicting self-frustrating oughts*

4.1. STRUCTURAL AMBIGUITY + DIVERSITY – QUASI-MT

Can we avoid the problem by giving up the validity of quasi-MT? We might do so by allowing the ordering for epistemic *must* to be nontrivial. Doing so gives the epistemic modal the same logic as the deontic modal. *must p* will be compatible with *not p*, and so quasi-MT will be no longer valid.

But for the purposes of the Death in Damascus example, this strategy doesn't help: we can simply stipulate that the relevant agent is absolutely certain that Death is a perfect predictor of his whereabouts. And so we can stipulate that it's a feature of the case that in every epistemically possible world where the agent goes to Aleppo, Death goes too – and so if he goes to Aleppo, it would be better for him to stay in Damascus.

Alternatively, it might be thought that we can reject the claim that in this circumstance, there is a quasi-MT violation. It is sometimes thought that the Death in Damascus case generates a rational dilemma, and so (25c) is false.

For every location you could go to, Death will have predicted that you'd go there. So we can construct similar conditionals for all such locations: *if you go to x, you should not go to x*.

There are two problems. First, even if this assessment of the rational norms is correct, it doesn't help us preserve the standard Kratzer semantics. That account rules out dilemmas: the modal background must be nonempty, and therefore the domain must be nonempty.

Second, if we generalize Kratzer semantics to allow for dilemmas, it's still a cost to a theory to predict that there *must* be a dilemma here. Causal decision theory doesn't treat this as a dilemma. In Section 3.4 I argued that the conditionals should be able to express verdicts of causal decision theory. So even if we wanted to accommodate dilemmas, there's reason to allow that the Death in Damascus case needn't generate a dilemma.

JOBNAME: No Job Name PAGE: 19 SESS: 66 OUTPUT: Mon Oct 6 17:21:13 2014 SUM: 46D7226A
/Xpp84/wiley_journal_P/PAPQ/papq_v0_i0/papq_12048

Finally, even if we embrace a theory of modals that allows for deontic dilemmas, we still need an explanation of the truth of the conditional. So far we have none. As I'll show in Section 4.2, the classical account faces challenges with predicting, and even permitting, the truth of the conditional.

4.2. PREDICTING THE CONDITIONAL

How do we predict the truth of a conditional like *If you go to Aleppo, you should not go to Aleppo*? Obviously, the single-modal reading is not an option with the restrictor semantics: it's predicted to be trivially false. So let's consider how we might find the double-modal reading. What values for the contextual parameters can plausibly be projected from context to predict the conditional?

Let A and D represent the propositions that Death goes to Aleppo and that Death goes to Damascus, respectively; and let a and d represent the propositions that you go to Aleppo and Damascus, respectively. Again, the deontic domain is the area surrounded by the dashed outline.

1. Suppose, naturally, that the deontic modal's modal background contains all four possibilities for where you and Death go: {Ad, Da, Aa, Dd}. (See Figure 4.)

Then the domain will include both kinds of worlds where you avoid Death. Among these will be worlds where you go to Damascus and worlds where you go to Aleppo. So it's not the case that at that modal background you should not go to Aleppo (going is permissible). *If you go to Aleppo, you shouldn't* is therefore predicted to be false.

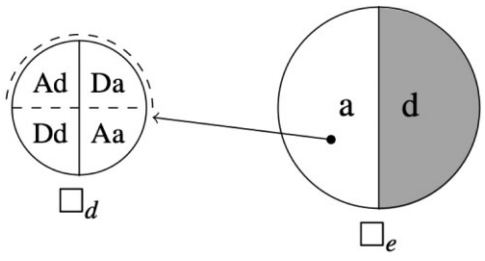

Figure 4

JOBNAME: No Job Name PAGE: 20 SESS: 66 OUTPUT: Mon Oct 6 17:21:13 2014 SUM: 5037DADD
/Xpp84/wiley_journal_P/PAPQ/papq_v0_i0/papq_12048

2. Suppose the restriction from the *if*-clause is inherited by both the epistemic and the deontic modal. The deontic modal's modal background includes only {Aa, Da}. (See Figure 5.)

Then the domain includes only worlds where you go to Aleppo. So the conditional *if you go to Aleppo, you shouldn't* is predicted to be false; indeed, on this reading, *if you go to Aleppo, you should* is predicted to be true.

3. Suppose the deontic modal's modal background, by (ad hoc?) stipulation, contains only worlds where Death goes to Aleppo: {Aa, Ad}. (See Figure 6.) Then the best worlds are worlds where you don't go to Aleppo. So we can get the prediction that *If you go to Aleppo, you shouldn't* is true.

There are two problems with this way of fixing the contextual parameters. First: what is the justification for holding Death's location fixed throughout the modal background? Presumably the information that Death will be in Aleppo is an inference from (a) the information from the antecedent that you will go to Aleppo and (b) the background information that Death will be where you are. But while the modal background reflects the conclusion from these two premises, it doesn't reflect the premises

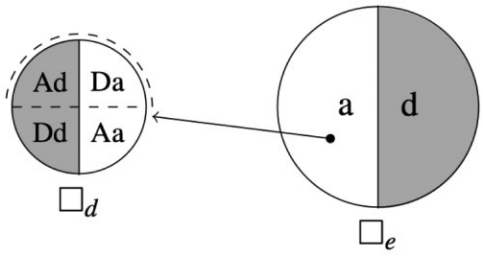

Figure 5

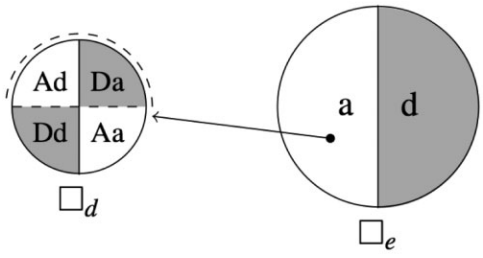

Figure 6

JOBNAME: No Job Name PAGE: 21 SESS: 66 OUTPUT: Mon Oct 6 17:21:13 2014 SUM: 473C9519
/Xpp84/wiley_journal_P/PAPQ/papq_v0_i0/papq_12048

themselves: the modal background doesn't reflect premise (b). It allows that you might be in a different place from Death, since it includes both Aa-worlds and Ad-worlds.

In other words, speakers must assume *in the very same breath* that Death must be in the same place as you and that he might not be. How is it possible for context to fix the sets of relevant circumstances in this way? It's not just that the stipulation of the modal background is ad hoc: it's not clear that this is even a coherent configuration of background assumptions.

A second and graver problem: there is a reading of (25a) and (25c) according to which both can be true. Let (25a) and (25c) be true at a point of evaluation $\langle w, f, g\rangle$. Then the deontic modal's domain at $\langle w, f, g\rangle$ contains both a-worlds and d-worlds.

Now assume (25a)'s modal background and domain are as represented in Figure 6. Then relative to $\langle w, f, g\rangle$, every epistemically possible a-point has a deontic domain that includes only d-worlds (since the domain in Figure 6 entails Ad). Therefore at all epistemically possible a-worlds, the deontic domain does not equal the deontic domain of $\langle w, f, g\rangle$. Therefore the actual point of evaluation can't be an a-point.

But we could run the same with (25b):

(25b) If you stay in Damascus, you shouldn't stay in Damascus.

So all the epistemically possible points will have a different deontic domain from the deontic domain at the actual point of evaluation. It follows that the world of evaluation can't be among the epistemically possible worlds. In other words, from the point of view of the speaker, the actual world won't be epistemically possible.

This problem will afflict any stipulation of modal background and ordering source we could use to predict (25a). And so with this kind of semantics for iffy *ought*s, there's no good way to predict the mutual truth of the sentences in (25) without treating them as equivocal.

But these sentences should be compatible at a single context. They are all the deliverances of a unified and coherent body of norms: namely, causal decision theory plus the desire to avoid death.

5. *A hypothesis*

To account for the consistency of the sentences in (25), it's helpful to see how other formal systems predict their mutual consistency.

First, note that the evidential expected utility of not going to Aleppo, given that you go to Aleppo, is not defined. It would have to be calculated

with $Pr(A \mid a \wedge \neg a)$, which is undefined because $Pr(a \wedge \neg a)$ is 0. But the causal expected utility of not going to Aleppo, given that you go to Aleppo, is defined.

For simplicity, we'll use a simple causal decision theory that uses (non-backtracking) counterfactuals: the causal expected utility of an act X is $Pr(X \mathbin{\square\!\!\rightarrow} S)U(XS)$ for all possible states S.

Suppose $Pr(A \mid a) = Pr(D \mid d) = .99$. Assume, as the example stated, that your acts do not cause any of Death's acts; he's merely a very good predictor of your actions. So $Pr(a \mathbin{\square\!\!\rightarrow} \mathrm{A}) = Pr(\mathrm{d} \mathbin{\square\!\!\rightarrow} \mathrm{A}) = Pr(\mathrm{A})$ and $Pr(a \mathbin{\square\!\!\rightarrow} D) = Pr(d \mathbin{\square\!\!\rightarrow} D) = Pr(D)$. The decision problem simplifies to:

	D	A
d	−100	0
a	0	−100

We are interested, not in each act's expected utility *simpliciter*, but in its expected utilities conditional on your going to Aleppo.

$$\begin{aligned} \mathrm{EU(d \mid a)} &= \mathrm{Pr(d \mathbin{\square\!\!\rightarrow} A \mid a)U(Ad) + Pr(d \mathbin{\square\!\!\rightarrow} D \mid a)U(Dd)} \\ &= -1 \end{aligned}$$

$$\begin{aligned} \mathrm{EU(a \mid a)} &= \mathrm{Pr(a \mathbin{\square\!\!\rightarrow} A \mid a)U(Aa) + Pr(a \mathbin{\square\!\!\rightarrow} D \mid a)U(Da)} \\ &= -99 \end{aligned}$$

So given that you go to Aleppo, not going to Aleppo has much higher causal expected utility. That's why (25a) can be used to express a result of causal decision theory:

(25a) If you go to Aleppo, you should not go to Aleppo.

Causal utility theory can be used to calculate the expected utility of an act conditional on your performing another, incompatible act because it separates two roles for acts. One role is the efficiently producing an outcome. Another is simply as information about the how the world is: it's a world where a particular act is performed. Just as we can calculate the expected utility of going to Aleppo conditional on its being sunny, so we can calculate the expected utility of going to Aleppo conditional on your not going to Aleppo.

If we want to allow that (25a), (25b), and (25c) can all be true without equivocation between them – that is, without changes in the contextually salient features, such as information and priorities – then I think we should generalize this lesson. We can separate the two different roles of the information in the *if*-clause:

JOBNAME: No Job Name PAGE: 23 SESS: 66 OUTPUT: Mon Oct 6 17:21:13 2014 SUM: 54A794F1
/Xpp84/wiley_journal_P/PAPQ/papq_v0_i0/papq_12048

1. the role of (at least potentially) changing which actions are available, and;
2. the role of (at least potentially) changing which actions are better than which.

Restricting the modal background takes care of the first role, while updating the probability function takes care of the second.

Kratzer semantics for modals says that information can affect what possibilities are available, by restricting the modal background. More recently, Kolodny and MacFarlane (2010) and others[25] have argued that information can affect not only what possibilities are available, but also which possibilities are better than which, by affecting the deontic ordering.

Examples like the Death in Damascus case show that information can affect availability and ideality independently of each other. In other words, information can change which relevant worlds are better without changing which worlds are available (and, of course, vice versa). They also show that we need these distinct roles represented distinctly in the semantics for deontic modals – i.e. two different informational parameters.[26]

The minimal generalization of Kratzer semantics that accomplishes this makes the ordering g a function of an information parameter. Instead of g providing an ordering, g is a function from a deontic information input to an ordering. In addition to the informational background f, which provides information about which worlds are available, we add a deontic information input parameter i, which is an input to g. Kratzer semantics is recoverable within this framework: wherever Kratzer semantics correctly predicts a particular ordering, the framework allows g to be a constant function from i to the relevant ordering.

DOMAINS: $domain(w, f, g, i)$ is the set of worlds in the modal background f ranked highest by the ordering $g(i)$.
MODALS: *ought p* is true at $\langle w, f, g, i\rangle$ iff p is true at all $w' \in domain(w, f, g, i)$.

The simplest assumption is that, with unrestricted modals, $f = i$. But we need not make this assumption. f might represent circumstantial information, e.g. the set of worlds compatible with what the agent can do. i might represent epistemic possibilities. We can remain neutral about how to model i: it might be a set of worlds, or a probability function, or something else.

In conditionals, the revision is more interesting. There are two basic strategies we might pursue for putting our extra parameter to work in cases like Death in Damascus.

1. Single-modal hypothesis: the conditional only has a deontic modal. The *if*-clause updates the i parameter that provides the input to the

prioritizing parameter *g*. But the antecedent doesn't necessarily update the *f*. That is, in some cases it doesn't as a restrictor for a quantifier.

So in (25a), the information that you'll go to Aleppo updates *i*, the information that affects what's best. This information alters the ordering such that worlds where you don't go to Aleppo are highest ranked within the unrestricted modal background *f*.

This strategy would generate an alternative to the restrictor analysis of conditionals. Benefits: we can explain self-frustrating iffy *ought*s while retaining quasi-MT (and, if we want to, DIVERSITY). We also don't make our semantics rule out expressing the norms of causal decision theory. Costs: we have to give up the restrictor analysis of conditionals, at least in its full generality. We also have to generate a means of predicting when antecedent information will restrict the modal background *f* and when it will not. These costs are substantial.

2. Double-modal hypothesis: A very similar option that retains the restrictor analysis: the *if*-clause on a self-frustrating iffy *ought* restricts the modal background of covert modal but not the overt modal embedded under it; but the *if*-clause also updates *i* for the deontic modal.

This strategy leaves us with all the benefits of the single-modal hypothesis with one exception: quasi-MT is no longer valid.

Of the two options, the second is more conservative and doesn't require updating the traditional story about how to predict ordinary single-modal iffy *ought*s. That is a point in its favor.

Let's take on board the assumption of a single-/double-modal ambiguity. The single-modal case will be probably be uninteresting: the information from the antecedent will update the both *f* and *i*. (What precisely this update amounts to in the case of *i* is left unspecified: if it's a set of worlds, presumably set intersection; if a probability function, it could be conditionalization, imaging, etc.)[27]

CONDITIONALS (overt modal restricted):
if p, ought q is true at $\langle w, f, g, i\rangle$ iff *q* is true at domain$(w, f + p, g, i + p)$

The double-modal case is more complex. In order to include modal backgrounds and orderings for both the covert epistemic modal and the overt deontic modal, the contextual parameters have to get more complex. I subscript deontic parameters with *d* and epistemic parameters with *e*.

JOBNAME: No Job Name PAGE: 25 SESS: 66 OUTPUT: Mon Oct 6 17:21:13 2014 SUM: 53C5558A
/Xpp84/wiley_journal_P/PAPQ/papq_v0_i0/papq_12048

Here the *if*-clause information plays two roles:

1. It restricts the epistemic modal's modal background f_e.
2. It updates the information input i for the deontic modal. (It does not restrict the deontic modal's modal background.)

CONDITIONALS (covert modal restricted):
if p, ($must_e$) $ought_d$ q is true at $\langle w, f_e, f_d, g_e, g_d, i\rangle$ iff for all $w \in$ *domain*$(w, f_e + p, g_e, i + p)$, for all $w' \in$ *domain*$(w', f_d, g_d, i + p)$, q is true.

So, for example, in the case of (25a), we look at all epistemically possible a-worlds, and check whether at each, the highest $g(i + a)$-ranked worlds that are compatible with the agent's capacities (f_d, including both a-worlds and d-worlds) are all worlds where the agent doesn't go to Aleppo. $g(i + a)$ might be, for example, a ranking of worlds in terms of causal expected utility, whether the expectation is determined by salient information updated with the information that the agent goes to Aleppo. The highest ranked worlds relative to these parameters, as we saw, were all worlds where the agent does not go to Aleppo.[28]

This provides a minimal generalization of Kratzer semantics plus the single-/double modal ambiguity that allows the two roles for information in deontic modals to vary independently of each other. Note that while this semantics for the double-modal conditionals allows the deontic modal to be sensitive to the antecedent's information – because i is updated with the antecedent's information – this need not make a difference. In many contexts, the salient priorities are not information-sensitive.[29]

6. Conclusion

I have argued that the *if p, ought p* puzzle causes deeper problems than previously realized. Kratzer semantics for conditionals in its current form can't predict or model the data. What the *if p, ought p* problem reveals is more general. First, there is an ambiguity between two forms of iffy *ought*: those that express conditional obligations and those that express speculation about obligation. Second, deontic modals are information-sensitive in two different ways: information can affect what possibilities are available and how these possibilities are ranked. These two forms of information-sensitivity can vary independently of each other. I have provided a generalized possible worlds semantics that is flexible enough to accommodate both of these facts.

School of Philosophy, Religion, and History of Science
University of Leeds

NOTES

[1] Kratzer, 1979, 1981, 1991, generalized from Lewis, 1975.

[2] To avoid typesetting complications, I use italics as a notational variant of corner quotes.

[3] This is meant to be an intuitive characterization of the two moving parts; I'm not committing myself to a particular view about deontic modals needing a particular flavor of modal background. The relevant circumstances may e.g. be epistemic.

[4] Similarly for other necessity modals: *must, have to, should,* etc. The classic account doesn't distinguish between so-called 'strong' versus 'weak' necessity modals. While I suspect there might be complex interactions between the strong/weak necessity modal distinction and the linguistic phenomena I discuss, they fall outside the scope of this article.

[5] A sociological note: throughout this article I use 'ought' as my canonical deontic necessity modal. Among philosophers ought is commonly used in expressions of moral, prudential, and rational necessity. ('Ought implies can'; 'subjective versus objective ought'; 'deriving an ought from an is'.) Nowadays 'ought' is typically not as natural in English outside philosophy, and so the examples I use may be better expressed with 'should'.

[6] The preorder, in Kratzer semantics, is projected as follows: $w > w'$ – w is 'more ideal' than w' – iff the set of ordering source propositions that w satisfies is a proper superset of the set of ordering source propositions that w' satisfies. Failures of trichotomy are irrelevant for my discussion, so for expository and diagrammatic purposes, total orders are more convenient.

[7] My examples use epistemic must because epistemic ought introduces a confound: the conditionals are sometimes judged infelicitous because, plausibly, epistemic ought generates a scalar implicature that its prejacents epistemic standing is somehow shaky. The trivial reading is easier to hear when the utterances include a sarcastic tone:

(i) a. If it's raining, it ought to be raining.
b. If it's raining, it's probably raining.

[8] Zvolenszky (2006) argues that the epistemic case is still problematic in some cases, with the following example:

(i) a. If this is the M3 motorway, then it must be the M3 motorway.
b. That is, If this is the M3, then I know that it is the M3.

If I'm interpreting Zvolenszky correctly, she claims that (ia) is not trivial, on the grounds that it can be glossed with (ib), which is obviously not trivial. I'm unconvinced that (ib) can gloss (ia) – I can't detect any such reading. Certainly this is not generally the case when the contextually salient epistemic possibilities are compatible with the speaker's knowledge. For example:

(ii) a. If the lights are out, she must not be home.
b. *Not equivalent to:* If the lights are out, then I know that she's not home.

See also further (more complex) evidence in Yalcin, 2007.

[9] Thanks to Stephen Yablo for discussion.

[10] Qualification: there are counterexamples where the antecedent intersected with the modal background is empty; but otherwise this is universally true, and there are some reasons to think that natural language modals are never evaluated at empty modal backgrounds. See discussion of the diversity condition below.

[11] Of course, this could turn out to be a case of so-called non-catastrophic presupposition failure, where sentences with presupposition failure are judged to be truth-evaluable: for example, 'The King of France is sitting on my lap.' See Yablo, 2006.

[12] It can also be generic, e.g. *If Ryan gets home early, he drinks before dinner*.

[13] Alternatively, we might predict that iffy oughts are always doubly modalized, but the *if*-clause can sometimes restrict both of the modals simultaneously. There is then an open question of how to predict when the *if*-clause restricts only one. I'll mostly ignore this alternative, but it's easy enough to generalize my discussion of the structural ambiguity account to this variant model.

[14] Note that the doubly modalized reading should sometimes be available even if the singly modalized reading is also available, as in Moss, 2005 manuscript, cited in von Fintel, 2011. §

[15] Mark Schroeder (pc) suggests that instead of suggesting that there's no trivial reading of *if p, ought p,* we can explain its apparent inaccessibility as an effect of its triviality. This may be so. It's worth noting, though, that trivial readings don't usually disappear just because a sentence has non-trivial disambiguations. For example: *If there are cranes on the esplanade, then there are cranes on the esplanade.*

[16] Kolodny and MacFarlane, 2010; Cariani, Kaufmann and Kaufmann, 2013; Charlow, 2013; Carr, 2012; Silk, 2014.

[17] Kolodny and MacFarlane attribute the case to Parfit (unpublished).

[18] For defenses, again, see Kolodny and MacFarlane, 2010; Cariani, Kaufmann and Kaufmann, 2013; Charlow, 2013; Carr, 2012; and Silk, 2014. For dissident voices, see Dowell, 2013; Bronfmann and Dowell, forthcoming.

[19] Assuming that epistemic modals are universal quantifiers over epistemically possible worlds (i.e. that their ordering is trivial).

[20] Kratzer semantics doesn't validate modus tollens, and counterexamples in natural language are easy enough to find. For example, gentle murderer cases: suppose at world w, S will murder. The speaker need not know this. Suppose the contextually salient modal background includes both murder- and not-murder-worlds, and all best worlds are not-murder-worlds. When the modal background is restricted to the murder-worlds, the best worlds in the restricted base are gentle-murder-worlds. Kratzer semantics says that at w, relative to contextual parameters satisfying these constraints, the following three sentences are all true:

(i) a. S ought not murder gently (or otherwise).
b. If S will murder, S ought to murder gently.
c. S will murder.

See e.g. Kolodny and MacFarlane, 2010 and Charlow, 2013. Note: if the speaker knows S will murder, then in order for sentences to be consistent within Kratzer semantics, the modal background can't be the set of worlds compatible with the speaker's knowledge.

[21] By 'respect', I mean preserve acceptance: if the premises are accepted at a context, then the conclusion must be accepted. (See e.g. Yalcin, 2007.) I use this notion as a means of retaining neutrality between different conceptions of validity: acceptance preservation versus truth preservation.

- p is accepted in an information state i iff $\forall w \in i : [[p]]^{w,i} = 1$.
- Consequence = acceptance preservation: $p_1,\ldots, p_n \vDash q$ iff every i that accepts $p_1, \ldots, p_n$ accepts q.
- $[[\Box_e q]]^{w,i}=1$ iff i accepts q.
- $[[\textit{if } p, \Box_e q]]^{w,i} = 1$ iff $i \cap [[p]]^{w,i}$ accepts q.

So $\neg q$, (*if* p, $\Box_e q$) $\models \neg p$. And then $\neg\Box_d q$, (*if* p, $\Box_e\Box_d q$) $\models \neg p$. But this holds because of epistemic modality's special relation to acceptance preservation; the deontic modal has no such relation.

[22] Examples like this are discussed in a decision theoretic context in Gibbard and Harper, 1978 and Hare and Hedden, 2012.

[23] Allan Gibbard and William Harper (1978) discuss this story, putatively an ancient Mesopotamian myth, in the context of decision theory; they cite Somerset Maugham's *Sheppey* (1934).

[24] Throughout I'll assume that the relevant kind of causal decision theory allows predictions of this form even when the antecedent concerns the agent's acts.

[25] Cariani, Kaufmann and Kaufmann, 2013; Charlow, 2013; Carr, 2012; Silk, 2014.

[26] I briefly defend a view where there are two informational parameters in Carr, 2012.

[27] It's worth exploring whether it's worthwhile to allow that in the single-modal case, i doesn't update with the antecedent information, but I won't pursue this here.

[28] Objection: If we're assessing actions with probability zero, why not assess the value of causing a miracle and killing Death? Reply: That option is not available at all in the modal background. That's why the modal background and deontic informational input are separate. This isn't stipulative: part of the purpose of a circumstantial modal background on a deontic modal is to set aside actions that an agent is not able to perform. But in many circumstances, even if you will p, you are able to not-p. (The fact that I will walk home tonight doesn't mean that I'm incapable of not walking home tonight.) So an antecedent *you* p shouldn't remove from the deontic modal background all $\neg p$ possibilities. It might be that I will do things I shouldn't.

[29] Distinguish information-sensitivity in contextual parameters from information-sensitivity in norms. A non-information-sensitive priority parameter g is a constant function from i to an ordering. Such a parameter might fully well capture the ranking of worlds provided by information-sensitive norms. For example, there can be a constant function from i to worlds where agents maximize expected utility relative to their own information, or the speaker's information (rather than the contextually salient information, i).

REFERENCES

Bronfman, A. and Dowell, J. (forthcoming). 'Contextualism about Deontic Conditionals,' in N. Charlow and M. Chrisman (eds) *Deontic Modality*. Oxford: Oxford University Press.

Cariani, F., Kaufmann, M. and Kaufmann, S. (2013). 'Deliberative Modality under Epistemic Uncertainty', *Linguistics and Philosophy*, 36, pp. 225–259.

Carr, J. (2012). 'Subjective "Ought".' Manuscript.

Charlow, N. (2013). 'What We Know and What to Do.' *Synthese* 190(12), pp. 2291–2323.

Dowell, J. (2013). 'Flexible Contextualism about Deontic Modals: A Puzzle about Information-sensitivity,' Special issue of *Inquiry* devoted to Contextualism and Relativism (H. Cappelen and S. Shapiro, eds), 56(2–3), pp.149–178.

Frank, A. (1997). *Context Dependence in Modal Constructions*. PhD thesis, Universität Stuttgart.

van Fraassen, B. C. (1972). 'The Logic of Conditional Obligation,' *Journal of Philosophical Logic* 1(3–4), pp. 417–438.

von Fintel, K. (2011). 'Conditionals,' in C. Maienborn, K. von Heusinger and P. Portner (eds) *Semantics: An International Handbook of Natural Language Meaning*. Berlin: De Gruyter Mouton.

Geurts, B. (2004). 'On an Ambiguity in Quantified Conditionals.' Manuscript.

Gibbard, A. and Harper, W. (1978). 'Counterfactuals and Two Kinds of Expected Utility,' in C. A. Hooker, J. J. Leach and E. F. McClennan (eds) *Foundations and Applications of Decision Theory*. Dordrecht: D. Reidel, pp. 125–162.

Hare, C. and Hedden, B. (2012). 'Self-Reinforcing and Self-Frustrating Decisions.' Manuscript.

Jackson, F. (1985). 'On the Semantics and Logic of Obligation,' *Mind* 94(374), pp. 177–195.

Kolodny, N. and MacFarlane, J. (2010). 'Ifs and Oughts,' *Journal of Philosophy* 107(3), pp. 115–143.

Kratzer, A. (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, A. (1991). 'Modality,' in A. von Stechow and D. Wunderlich (eds) *Semantics: An International Handbook of Contemporary Research*. Berlin: de Gruyter, pp. 639–650.

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

Maugham, S. (1934). *Sheppey*. New York: Doubleday.

Parfit, D. (unpublished). 'What We Together Do.'

Prior, A. N. (1958). 'Escapism: The Logical Basis of Ethics,' in A. I. Melden (ed.) *Essays in Moral Philosophy*. Seattle, WA: University of Washington Press, pp. 135–146.

Silk, A. (2014). 'Evidence-Sensitivity in Weak Necessity Deontic Modals,' *Journal of Philosophical Logic* 43(4), pp. 691–723.

Spohn, W. (1975). 'An Analysis of Hansson's Dyadic Deontic Logic,' *Journal of Philosophical Logic* 4(2), pp. 237–252.

Yablo, S. (2006). 'Non-Catastrophic Presupposition Failure,' in J. Thomson and A. Byrne (eds) *Themes from the Philosophy of Robert Stalnaker*. Oxford: Oxford University Press, pp. 164–190.

Yalcin, S. (2007). 'Epistemic Modals,' *Mind* 116(464), pp. 983–1026.

Zvolenszky, Z. (2002). 'Is a Possible-Worlds Semantics of Modality Possible? A Problem for Kratzer's Semantics,' in B. Jackson (ed.) *Proceedings of Semantics and Linguistic Theory 12*. Ithaca, NY: Cornell University Press, pp. 339–358

Zvolenszky, Z. (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). Budapest: Eotvos Lorand University Press, pp. 167–177.

Zvolenszky, Z. (2007). 'The Lost Pillar of Deontic Modality,' *Modality, Names and Descriptions*, PhD thesis, New York University.