

The best we can (expect to) get?

Challenges to the classic semantics for deontic modals*

Kai von Fintel (MIT), January 9, 2012

Paper for a session on Deontic Modals at the Central APA, February 17, 2012

There has been a recent upswing of work on deontic modality that has tried to overthrow what might reasonably be called the classic analysis of deontic modality in natural language. Here, we are going to look at a subset of the complaints lodged against the classic semantics and try to explore how to withstand the assault. We may not establish once and for all that the classic approach is the best approach we can get, but I hope to show that it is not that bad after all.

1 Introduction

1.1 The classic semantics for deontic modals

The classic semantics for modals has two contextual components, a *modal base* of accessible worlds and an *ordering* of the worlds in the modal base.¹

* Thanks to Thony Gillies and Angelika Kratzer for collaborating on this topic. Some of the themes in this paper were touched on in a talk “Deontic logic and natural language” (with Thony Gillies) given in a special session on “Logic and Linguistics” organized by the Association for Symbolic Logic during the 85th Annual Meeting of the American Philosophical Association Pacific Division. San Diego, CA. April 21, 2011. Exchanges with Fabrizio Cariani, Nate Charlow, Janice Dowell, Magda Kaufmann, Stefan Kaufmann, and Alex Silk (some of these in communications with Angelika Kratzer) were helpful. There is some overlap with the discussion of (non-)monotonicity of desiderative attitudes in Luka Crnic’s dissertation (2011).

¹ The classic semantics for deontic modality was developed in variants by Danielsson 1968, Hansson 1969, Føllesdal & Hilpinen 1971, van Fraassen 1972, Lewis 1973, 1974, and von Kutschera 1974 (an incomplete listing of early work). It shares some properties with an even more ancient and bare bones deontic logic system, which is often called SDL (Standard Deontic Logic). In deontic logic circles, the classic system is sometimes called Dyadic Deontic Logic (because of the assumed logical syntax of giving the obligation operator two arguments standing for the given circumstances and the prejacent.). It was brought into linguistic semantics by Kratzer 1977 (earlier German versions were published in 1975 and 1976), Kratzer

Modals quantify over the best worlds in the modal base. In the case of deontic modality, the modal base is typically a set of *circumstantially accessible* worlds, that is worlds that match the evaluation world as far as a certain set of circumstances are concerned, circumstances that are considered immutable, not (anymore) in the control of those whose obligations/permissible options are the topic at hand. The worlds in the modal base are then ordered by how they match the ideals/precepts of a body of law, ethics, requirements. Deontic modals then quantify over the best worlds in the modal base, the worlds that given the circumstances are ordered as best among the remaining possibilities. Deontic modals like *ought* or *have to* are necessity modals that claim that their prejacent is true in all of the best worlds. The corresponding possibility modal is *may*, which as the dual of the necessity modals claims that the prejacent is true in some of the best worlds.

We said above that the two contextual components of deontic modals are a set of accessible worlds and an ordering of those worlds. Because deontic modals can themselves occur in intensional contexts, the contextual parameters will actually need to be functions from evaluation worlds to accessible worlds and to orderings:

- (1) $\llbracket \text{OUGHT } \phi \rrbracket^{MB, \geq} = \lambda w. \forall w' \in \text{BEST}_{MB, \geq}(w): \phi(w') = 1.$
 where $\text{BEST}_{MB, \geq}(w) = \{w' \in MB(w): \neg \exists w'' \in MB(w): w'' >_w w'\}$

A variant, due to Kratzer, is to think of the contextual components as (functions from evaluation worlds to) sets of propositions: one set circumscribes the circumstances (and taken together identifies the worlds in the modal base as those in which all the relevant circumstantial propositions are true) and the other set is used to evaluate and order the worlds in the modal base, counting one world as better than or equally as good as another if any proposition in the set that is true in the latter is also true in the former. These (functions from evaluation worlds to) sets of propositions are known as *conversational backgrounds* in the Kratzerian framework. The conversational background that induces the ordering on worlds is called the *ordering source*.

- (2) $\llbracket \text{OUGHT } \phi \rrbracket^{MB, OS} = \lambda w. \forall w' \in \text{BEST}_{MB, OS}(w): \phi(w') = 1.$
 where $\text{BEST}_{MB, OS}(w) = \{w' \in \bigcap(MB(w)): \neg \exists w'' \in \bigcap(MB(w)): \exists p \in OS(w): p(w'') = 1 \ \& \ p(w') = 0\}$

1981, and the influential handbook article Kratzer 1991. In what follows, we will often use terminology (such as *modal base*) established in linguistics but perhaps not as entrenched in adjacent disciplines.

The best we can (expect to) get?

With Kratzer (and Lewis), we might consider giving up on the Limit Assumption, which is behind the locution *the best worlds*. If so, the semantics would need to be formulated in a more complicated way. This consideration does not interact closely with the concerns of this paper, so we will make the Limit Assumption in what follows.

Kratzer famously advocated a treatment of conditionals that analyzes *if*-clauses as restrictors of (modal) operators. In particular, conditionals with deontic modals in the consequent are interpreted as the *if*-clause restricting the modal, that is, the modal base is narrowed to only those worlds in it that make the antecedent proposition true. While this makes good on some of the original motivations for the classic semantics (which grew at least partially out of attempts to solve various semantic puzzles surrounding conditional obligation), we will have little to say about conditionals until the very end.

It is important to realize that the study of natural language deontic modality is not simply the theory of the English word *ought*. Even in the case of English, we need to at least add *must*, *have to*, *should* to the equation, and of course, expressions of permission such as *may*. One recent concern that we will not address further is the distinction in strength between *must/have to* and *ought/should*, even though both are necessity modals. As argued in von Stechow & Iatridou 2008, this distinction may be due to the latter being sensitive to a more fine-grained ordering (perhaps modeled via a secondary ordering source). Another recent insight is that English *must* may prefer being used in a performative setting (where it doesn't just report on an obligation but actually puts it in force); see Ninan 2005.

1.2 Challenges

As mentioned, there have recently been challenges to the adequacy of the classic semantics. Two convenient venues where these complaints have been compiled are Cariani 2011 and Lassiter 2011 (see also Nauze 2008), but as will become apparent, the complaints mostly have a long history as do most of the obvious counter-moves available to the classic semantics.

There are five clusters of challenges: (i) problems having to do with right upward monotonicity (Ross' Paradox, Professor Procrastinate), (ii) moral dilemmas, (iii) information sensitivity (Miners Paradox), (iv) the interpretation of certain deontic conditionals (such as *if p, ought p*), (v) issues surrounding the (non-)gradability of deontic modal expressions.

Here, we will deal with the first three of these, leaving the other two aside.² The discussion here will remain admittedly programmatic. The point is to put down a challenge: the opponents of the classic semantics either overlook or too eagerly dismiss ways in which the classic semantics can account for the allegedly recalcitrant data. Further, in several areas, the proposed alternative semantics actually fail to do justice to the data.

2 Monotonicity

Cariani and, following him, Lassiter argue that the classic semantics is mistaken in its monotonicity. In the classic semantics with the Limit Assumption, **OUGHT** ϕ says that all of the best worlds in the modal base are ϕ -worlds and if ϕ -worlds are all ψ -worlds, all the best worlds will also be ψ -worlds and thus **OUGHT** ψ will be true automatically. So, for any ψ entailed by ϕ , **OUGHT** ϕ will entail **OUGHT** ψ .

We will review two of the cases adduced by Cariani that are meant to persuade us that claiming monotonicity is a bad idea. We argue that they are rather to be seen as two paradigm examples of well-known mechanisms creating the *appearance* of semantic non-monotonicity for natural language operators that are actually monotonic. Finally, we will argue that there are strong reasons to stick to the monotonicity of the classic semantics.

2.1 Ross' Paradox

2.1.1 The problem

Ross 1941 argued that the logic of imperatives should be non-monotonic, based on this example:

- (3) a. Slip the letter into the letter-box! $\not\Rightarrow$
 b. Slip the letter into the letter-box or burn it!

Someone who was given the command in (3a) would not be warranted in arguing that the command entailed the one in (3b) and perhaps then proceeding to burn the letter.

² The issues around *if p, ought p* are well-known in deontic logic and were raised in the linguistic literature by Frank 1996 and Zvolenszky 2002 and are discussed by Geurts 2004 and Kratzer 2010. The question of the (non-)gradability of modals, central to Lassiter's thesis, is investigated by Klecha 2011.

The best we can (expect to) get?

The argument transposes straightforwardly into examples with deontic modals (let's take the opportunity and adjust Ross' quaint language of slipping letters into letter-boxes):

- (4) a. You ought to mail the letter. $\not\Rightarrow$
b. You ought to mail the letter or burn it.

An inference from (4a) to (4b) doesn't seem valid. But surely, if all the best worlds are worlds where you mail the letter, then those are all worlds where at least one of these propositions are true: you mail the letter, you burn it. So, the classic semantics predicts the inference to be valid, incorrectly so. Before we turn to the debate on whether the absurdity of the inference reveals semantic non-monotonicity, we will show that the problem extends beyond the interaction of *ought* with disjunction.

2.1.2 Not just *ought*

The inference failure in (4) doesn't just arise with *ought*, it happens with all deontic necessity modals and actually, with deontic possibility modals as well:

- (5) a. You $\left\{ \begin{array}{l} \text{must} \\ \text{have to} \\ \text{need to} \\ \text{should} \\ \text{may} \end{array} \right\}$ mail the letter. $\not\Rightarrow$
b. You $\left\{ \begin{array}{l} \text{must} \\ \text{have to} \\ \text{need to} \\ \text{should} \\ \text{may} \end{array} \right\}$ mail the letter or burn it.

All the modals in (5) are treated as monotonic vis-à-vis their prejacent by the classic semantics and so the inference in (5) is incorrectly predicted to be fine in all cases.

2.1.3 Not just disjunction

Something like Ross' Paradox doesn't just arise with disjunction under deontic modals. There are other expressions that arguably can create similar

effects.

- (6) a. You ought to invite Professor Edgington. \nrightarrow
 b. You ought to invite some expert on conditionals.

Even if we assume as given that Professor Edgington is an expert on conditionals, and thus inviting Professor Edgington entails inviting an expert on conditionals, the inference in (6) is far from perfect.

The effect is more pronounced with expressions that are more aggressive markers of indifference, such as the German *irgendein* or the English *some ... or other* (which intriguingly, incorporates disjunction).

- (7) a. You ought to invite some expert on conditionals or other.
 b. Du solltest irgendeinen Experten über Konditionale einladen.

(6a) certainly does not warrant the inference to the conclusions in (7).

2.1.4 What needs to be explained

There is more to be explained than just the failure of these inferences. If that was all that was needed, the conclusion that deontic modals are simply non-monotonic in their semantics might be compelling, since it would explain why these inferences fail. But there are two other desiderata: (i) as we'll see later, there are compelling reasons to think that deontic modals are in fact monotonic and so we need an understanding of when monotonic inferences are safe and when they are not; (ii) there is a very clear intuition about *why* these inferences fail: the supposed conclusion itself licenses inferences that are not justified by the premise, and so we need to understand the grounds for those inferences from the conclusion. Let's focus on the second point for now. The deontic modal with a disjunction in its prejacent licenses an inference as to what is permitted:

- (8) a. You ought to mail the letter or burn it. \Rightarrow
 b. You may mail the letter.
 c. You may burn the letter.

Confronted with (8a), we infer both (8b) and (8c), that is we infer that we have free choice as to which of the two actions to take. It is quite obvious that it is this free choice inference that disrupts the problematic inference that constitutes Ross' Paradox. Saying that deontic modals are non-monotonic

The best we can (expect to) get?

cannot claim to be a full solution to Ross' Paradox unless it is accompanied by an explanation of how free choice inferences arise.

2.1.5 Approaches to free choice

There are two and a half approaches to free choice inferences that we'll mention here. One approach, taken by Aloni 2007 (see also Simons 2005), is to build the free choice inference into the semantics by making deontic modals (and other modal operators) sensitive to alternatives introduced by disjunction (and similar operators). A simple way to put it: **OUGHT** (ϕ or ψ) is true iff in any of the best worlds either ϕ or ψ is true and there are some best worlds where ϕ is true and some best worlds where ψ is true. At least one of the alternatives needs to be true in any of the best worlds and each alternative needs to be *represented* in some of the best of worlds. This semantics doesn't even take the prejacent to be a proposition (instead it takes a set of alternative propositions) and so the question of whether it is monotonic on the prejacent proposition is moot. It is, however, monotonic in an obvious sense when the prejacent presents a single proposition.

The other approach (and its younger half sibling) goes back at least to Hare 1967 (Cariani cites Wedgwood 2006 for this proposal). The idea is that the free choice inference is an *implicature* in a (Post-/Neo-)Gricean sense. We infer that either disjunct is a permitted option because if one of them wasn't then the speaker could have uttered the simpler sentence with just the other disjunct as the prejacent. That is, the reasoning takes flight from the fact that the speaker of (9a) chose not to say either (9b) or (9c):

- (9) a. You ought to mail the letter or burn it.
- b. You ought to mail the letter.
- c. You ought to burn the letter.

Exactly how to fill out the story is a bit intricate. For one, one has to justify why those simpler *ought*-sentences are relevant alternatives. We cannot go into the details here; cf. Chierchia, Fox & Spector 2012 for a general survey of implicature. That paper is also a good source for learning about a new perspective on how implicatures work, namely in a way that is not as *post-compositional* or pragmatic as assumed by (Neo-)Griceans but rather integrated into the recursive grammar.

Aloni-style approaches build the free choice inference into the semantics and thus quite fully explain the failure of Ross's Inference. The non-

monotonic approaches of Cariani and Lassiter appear to have remained silent on the status of free choice inferences, but since their systems are precisely designed to allow very few clearcut logical inferences, one assumes that they will have to join the classic semantics in treating free choice inferences as context-dependent, non-semantic inferences, most likely implicatures. (NB: Cariani and Lassiter do not comment on the source of free choice inferences, perhaps because they think that as soon as a semantics is non-monotonic, it needs to do no further work to account for Ross's Paradox. But this is simply not so.)

We believe that treating free choice inferences as implicatures, while leaving the semantics simple and monotonic, is the right approach. Far from being an instance of *explaining away* problematic data, this approach constitutes a bona-fide explanation of the data that we suspect will be more than competitive with non-monotonic approaches once those have been supplemented with a mechanism to explain free choice inferences.

2.1.6 Testing for implicature

One would have thought that it should be easy to test whether free choice inferences are indeed implicatures or whether they are part of the truth-conditional semantics of deontic modals (and thus providing an argument for an account like Aloni's).

Classically, we would test at least for the following properties of implicatures: cancelability, suspendability, reinforceability, tendency to vanish in embedded environments. Let's apply these to a textbook case: the inference from *some* to *some and not all*:

- (10) Cancelability
John ate some of the cookies. In fact, he ate all of them.
- (11) Suspendability
John ate some of the cookies, $\left\{ \begin{array}{l} \text{if not all of them} \\ \text{perhaps all of them} \end{array} \right\}$.
- (12) Reinforceability
John ate some of the cookies, but he didn't eat all of them.
- (13) Embedding
If John ate some of the cookies, he might get salmonella poisoning.

These are textbook applications of the tests to a textbook example of an

The best we can (expect to) get?

implicature. Nevertheless, the devil is, as always, in the details and much of what we thought we knew has come under critical scrutiny recently. So, it's not necessarily a straightforward matter to apply these tests to a case in dispute. Nevertheless, let's try, we might learn something. We'll use a disjunctive prejacent that Cariani employs: *wear a tie or a scarf*.

Cancelability

(14) John ought to wear a tie or a scarf. In fact, he ought to wear a tie.

As discussed by van Kuppevelt 1996, among others, *in fact* cancellations of implicatures are not all that simple. Some people (cf. Sauerland 2004) now think that certain implicatures cannot truly be cancelled at all, in the sense of being first generated and then withdrawn. The best kind of context for apparent cancellation of the *in fact* kind is one where the first sentence doesn't give rise to the implicature in the first place and the *in fact* phrase adds a strengthened conjunct. This seems right for our example as well.

(15) Q: Ought John to wear a tie or a scarf (to the interview)?

A: Hell, yeah, in fact he ought to wear a tie.

That seems fine.

Suspendability

(16) John ought to wear a tie or a scarf; and maybe, indeed, he ought to wear a tie.

Again, this doesn't roll off the tongue but doesn't seem fundamentally flawed.

Reinforceability

(17) John ought to wear a tie or a scarf, but the choice is up to him.

This is impeccable and argues that the free choice inference is something that is not intrinsic to *ought-or* but is a secondary inference that can be reinforced. The acceptability of (17) might present a problem for Aloni-style analyses.

Embedding In principle, as soon as an implicature-triggering sentence is embedded in an non-assertive context, a Gricean approach predicts that the implicature should vanish, since the sentence not being asserted means that we do not even start reasoning why the speaker asserted the sentence rather than one of its alternatives (of course, we do reason that way about the entire matrix sentence which includes our target sentence, but that is another matter). But things are not entirely straightforward. There are in fact plenty of cases where implicatures do not seem to vanish upon embedding. Both of the following still seem to carry their usual implicatures:

- (18) a. I believe that John ate some of the cookies.
 b. I believe that John ought to wear a tie or a scarf.

It is customary to use negative or, more generally, downward-entailing environments as tests, but even those bring issues. Some implicature-triggers resist embedding in negative contexts because they have positive polarity restrictions. *Some*, for example, doesn't like taking scope under negation (instead, its suppletive counterpart *any* is typically used in such environments). With deontic modals, the same issue arises. Both *must* and even more so, *ought* resist scoping under negation.³ As Lassiter discusses at some length, *ought* appears to be a neg-raiser:

- (19) a. I doubt that you ought to go there. \Rightarrow
 b. I think that you ought to not go there.

This means that it is advisable to use modals that do not resist scoping under negation. In particular, *have to* doesn't have any such compunction and should give rise to cleaner data. And indeed, it does:

- (20) a. I doubt that you have to wear a tie or a scarf.
 b. You don't have to wear a tie or a scarf.
 (21) a. Nobody has to wear a tie or a scarf. \Rightarrow
 b. Nobody has to wear a tie and nobody has to wear a scarf.

(20), where interfering free choice inferences are blocked by the embedding under *doubt*, means exactly what the classic semantics predicts: you neither have to wear a tie nor do you have to wear a scarf. It is very hard to imagine reading (20b) as denying free choice while maintaining that you have to

³ See Iatridou & Zeijlstra 2011 for some recent work on this topic.

The best we can (expect to) get?

wear one of the two neck decorations. Similarly, the inference in (21) seems impeccable and exactly what's predicted by the classic semantics. Alonso-Ovalle (2006: 141ff.) uses such data to argue for an implicature account and against a Simons/Aloni-style semantics for disjunctive prejacent. As he notes, the argument was first used by Kratzer & Shimoyama 2002 in their work on indeterminate pronouns in prejacent.

2.1.7 Cariani's Arguments

Cariani is well aware of the implicature response. But he believes he has a counter-argument: as we've seen, implicatures are expected to be cancelable, whereas it seems to him that this is not the case here. He compares the case of disjunction under *ought* with the case of disjunction under epistemic *must*. There are two pairs of examples that are supposed to show a clear difference.

- (22) a. <Cariani: #> I doubt that Lynn must have either worn a tie or a scarf. In fact, she must have worn a scarf.
b. I doubt that Lynn ought to either wear a tie or a scarf. In fact, she ought to wear a scarf.⁴

(23) The Bridge Game I

- A: Joe must have followed suit or played the king of trumps.
B: Jane had the king of trumps. He must have followed suit.
A: <Cariani: #> I guess what I said was wrong.

(24) The Bridge Game II

- A: (According to the rules), you ought to follow suit or play the king of trumps.
B: No, the rules quite explicitly say you ought to follow suit, no matter what.
A: I guess what I said was wrong.

Unfortunately for Cariani, these data points are less than persuasive. To get the second pair out of the way, there are two problems with it. First, the pair is far from minimal, so we should make them more parallel by changing the epistemic exchange as follows:

⁴ Lassiter 2011: 126 changes this to *I doubt that Lynn ought to either wear a tie or a scarf. In fact, I'm reasonably certain that she ought to wear a scarf.*

(25) The Bridge Game I'

A: Joe must have followed suit or played the king of trumps.

B: No, the evidence is quite clear: Jane had the king of trumps. He must have followed suit.

A: I guess what I said was wrong.

For what it is worth, we do not find (25) to be degraded in comparison to (24). Second, it is unclear why we should suppose that the *I guess what I said was wrong* response is a way to distinguish sharply between truth-conditional content and conversational implicature. Retraction/rejection phenomena can be sensitive to all kinds of dimensions of meaning and it seems they can easily target implicatures:

(26) A: Jane is in the kitchen or in the living room.

B: No way, she's never in the kitchen.

A: I guess what I said was wrong.

So, if anything, the problem here would be why in the epistemic case the implicature is more resistant to being used as a cause for rejection. But as we've said, we're not even sure whether the data point is correct.

Cariani's other data point combines embedding the modal with its disjunctive preajacent under *I doubt that* plus an *in fact* follow up, thus mixing two tests for implicature, embedding under a negative operator and an *in fact* cancellation. If things were simple we would expect that the first component simply negates the necessity claim and neither gives rise to free choice inferences nor negates them, and the second component would contradict the assertion of the first component by claiming the necessity of one of the disjuncts. Cariani claims that the epistemic variant behaves just as predicted:

(22a) <Cariani: #> I doubt that Lynn must have either worn a tie or a scarf. In fact, she must have worn a scarf.

But in the deontic case, Cariani reports no problem:

(22b) I doubt that Lynn ought to either wear a tie or a scarf. In fact, she ought to wear a scarf.

Or Lassiter's variant:

(27) I doubt that Lynn ought to either wear a tie or a scarf. In fact, I'm

The best we can (expect to) get?

reasonably certain that she ought to wear a scarf.

We have already seen that embedding *must* and even more so *ought to* under *doubt* or any negative operator is not easy because these modals do not like negative contexts. The epistemic case brings with it a further issue: embedding epistemic modals under operators like *doubt*, which are themselves epistemic in nature, raises questions about iterated epistemic modality, which we are loathe to get into here.

When we clean up the data by choosing a straightforward negation and the kind of modal that doesn't mind being negated, the data behave as expected from the point of view of classic semantics with free choice implicature:

- (28) a. Lynn doesn't have to have worn a tie or a scarf. # In fact, she has to have worn a tie.
b. Lynn doesn't have to wear a tie or a scarf. # In fact, she has to wear a scarf.

In neither case is the *in fact* follow-up allowed.

Our conclusion: there is every reason to think that the implicature account of what goes wrong in Ross' examples is on the right track. So, there is no reason to abandon the monotonic classic semantics.

2.1.8 We actually want monotonicity

In fact, in this discussion we have already seen strong reasons to maintain a monotonic classic semantics. Consider in particular the interpretation of negated deontic modals:

- (29) a. Lynn doesn't have to wear a tie or a scarf.
b. Nobody has to wear a tie or a scarf.

It seems clear that these sentences have straightforward monotonic entailments. (29a) entails that Lynn doesn't have to wear a tie. (29b) entails that nobody has to wear a tie. The following conjunctions are simply insane:

- (30) a. #Lynn doesn't have to wear a tie or a scarf, but, of course, she has to wear a tie.
b. #Nobody has to wear a tie or a scarf, but, of course, some people have to wear a tie.

The non-monotonic account cannot account for the fact that the negated modals have such a strong meaning and that hence, the conjunctions in (30) are so hopeless.

There are many more examples where these come from, even in the absence of disjunction or other alternative-offering expressions. Consider:

- (31) #You don't have to bring alcohol to the party, but you do have to bring wine.

On the assumption that wine is alcohol, (31) is not coherent, but a non-monotonic account would lead us to expect that this should be just as fine as the Ross-type examples.

Finally, consider the fact that negated deontic necessity modals license NPIs:

- (32) You don't have to bring any alcohol to the party.

Within the standard theory of NPI-Licensing, expressions like *any* require a downward-entailing environment. But if the non-monotonic semantics for deontic modals is correct, then *not have to* would not be downward monotonic (because *have to* is not upward monotonic). So, accounts like Cariani's or Lassiter's incur the cost of having to provide an alternative theory of NPI-licensing. We are not saying that this is a completely hopeless commitment, but we are content to stick with the monotonic classic semantics which considers (32) an unremarkable case of an NPI licensed by a downward monotonic environment.

2.2 Professor Procrastinate

Apart from Ross' Paradox, Cariani and Lassiter also proffer the famous case of Professor Procrastinate as an argument for a non-monotonic semantics for deontic modals. As the story goes (Jackson & Pargetter 1986):

Professor Procrastinate receives an invitation to review a book. He is the best person to do the review, has the time, and so on. The best thing that can happen is that he says yes, and then writes the review when the book arrives. However, suppose it is further the case that were Procrastinate to say yes, he would not in fact get around to writing the review. Not because of incapacity or outside interference or anything like that, but

The best we can (expect to) get?

because he would keep on putting the task off. (This has been known to happen.) Thus, although the best that can happen is for Procrastinate to say yes and then write, and he can do exactly this, what would in fact happen were he to say yes is that he would not write the review. Moreover, we may suppose, this latter is the worst that can happen. It would lead to the book not being reviewed at all, or at least to a review being seriously delayed.

Given this scenario, it seems clear that Procrastinate should not accept the review invitation, since it is given that he won't *in fact get around to writing the review* and thus would hurt the project. But if we were asked whether he should accept *and write* the review, we would surely say yes, that's what he should do. So, it seems that we consider (33) true but (34) false. That seems like a good reason to deny monotonicity.⁵

(33) Procrastinate ought to accept and write. \Rightarrow

(34) Procrastinate ought to accept.

Similar cases have been adduced by Asher 1987: p. 171 for bouletic attitudes. Consider Nicholas, who would of course not mind a free trip on the Concorde (not possible anymore, alas), but given that really there are no free trips on the Concorde, he doesn't actually want a trip on the Concorde. So, Asher concludes, the following is invalid:

(35) Nicholas hopes to get a free trip on the Concorde. So Nicholas hopes to get a trip on the Concorde.

We do not think that these kinds of cases argue for non-monotonicity. The counter-story has been told before, in the case of attitudes, in von Fintel 1999.⁶ The pattern in all of these examples is a (more or less) surreptitious context change. First, we are led to believe that the ideal course of action

⁵ Procrastinate is the most famous case, but similar scenarios are found earlier and elsewhere. For example, Kamp 1973: pp. 59-60: *Perhaps the man standing under the first floor window of the burning house ought, no matter what, to tell the frightened girl behind it to jump, and catch her in his arms when she does. But if he cannot catch her, or doesn't want to risk his bones, he certainly ought not to tell her to jump.*

⁶ The main problem with the argumentation by Asher was already identified by Kanazawa 1991. Asher had given the following as another invalid inference:

(i) Fred wants a car and a garage for it. So Fred wants a garage for a car.

(accepting and writing, a free trip on the Concorde) is not available, which then, under those realistic constraints, makes us assent to *Procrastinate ought not to accept* and *Nicholas doesn't want a trip on the Concorde*. But then, somehow, the possibility of the ideal is suddenly introduced as something not entirely unachievable. With respect to this new context, we of course assent to *Procrastinate ought to accept and write* and *Nicholas hopes to get a trip on the Concorde*. But the fact that the context has shifted cannot be ignored. With respect to the new context, we do not in fact anymore assent to the earlier assessments. It is not so that *Procrastinate ought not to accept*, what he should do is accept and write. We can see this in the flat-footed disaster of unadorned conjunction:

- (36) #*Procrastinate ought to accept and write, but he ought not to accept.*
 (37) #*Nicholas hopes to get a free trip on the Concorde, but he doesn't hope to get a trip on the Concorde.*

These are hopeless, but predicted to be just fine by a non-monotonic analysis.⁷

2.3 Conclusion

We have seen that there is good reason to think that deontic modals are *semantically* monotonic with respect to their prejacent, while there are at least two ways in which this semantic monotonicity can be obscured: implicatures triggered by salient alternatives (made salient by disjunctions or indifference indefinites) and context shifts. Such factors create a situation where even though modals are monotonic, upward inferences cast in natural language are not necessarily “safe”. This is quite reminiscent of what von Fintel 1999, 2001 and Gillies 2007 have discussed with respect to the left downward (non-)monotonicity of counterfactual conditionals. Champions of semantic

Kanazawa points out, thereby anticipating von Fintel's argument, that the following sounds like a downright contradiction:

- (ii) Fred doesn't want a garage, but he wants a car and a garage.

⁷ Asher seems to be slightly aware of the hopelessness of (37). In the gloss right below it, he says *If I want to ride on the Concorde and not pay for it, it doesn't necessarily follow that I also want to ride on the Concorde simpliciter*. The addition of *simpliciter*, or *no matter what*, or similar expressions, of course, strengthens the meaning of the *want*-claim to such an extent that it doesn't follow from the premise. But that is another form of sleight of hand, because the monotonic analysis would of course not claim an inference to such a strong conclusion.

The best we can (expect to) get?

non-monotonicity might do well to mine Moss 2010 for ways of resisting the conclusions we have reached here. In any case, the point is that the prima facie invalidity of certain inferences does not give direct evidence for *semantic* non-monotonicity. It is worth saying that semantic non-monotonicity is likely inevitable if one works with one's hands tied behind the back, as deontic logicians typically do, since they do not even try to fully model the complexity of natural language meaning with its multiple dimensions (truth-conditional semantics, presuppositions, dynamic contexts, conventional implicatures, conversational implicatures, etc.).

3 Moral dilemmas

I have a dilemma. My tax return is due by midnight tonight. So, I ought to work on that. But, my review for L&P is also due tonight. So, I ought to work on *that*. But I can't do both. So, I'm screwed, as happens far too often. What should I do?

That is a rather quotidian example of a moral dilemma. If you want fancy, there are much more dramatic ones (like Sartre's famous case of the wartime youth torn between joining the resistance and staying home to take care of his elderly mother) and even arbitrarily constructed ones — Horty 2003 imagines having simultaneously made identical and incompatible promises to two identical and identically situated twins. We can stick with the everyday here. What is important is that it does not seem obviously crazy to imagine that there might be conflicting obligations and that one might truthfully describe such a situation with **OUGHT** ϕ and **OUGHT** ψ where ϕ and ψ are incompatible (either intrinsically or given the facts on the ground).

Lassiter (2011) complains that the classic semantics for deontic modals has no space for such moral dilemmas and offers his own analysis as a remedy. We show in this section that in fact, the classic semantics does offer some space for moral dilemmas, albeit requiring some modifications, and that Lassiter's remedy is not a good one.

Without any further modification, the classic analysis doesn't allow dilemmas. **OUGHT** ϕ is true iff all the best worlds in the modal base are ϕ -worlds, and it can't be that all the best worlds are also ψ -worlds, if ϕ and ψ are incompatible, and so **OUGHT** ψ cannot be true at the same time as **OUGHT** ϕ . And even the more intricate constructions proposed in the absence of the limit assumption do not allow for incompatible prejacent. What the classic approach does is to *adjudicate* any apparent conflicts, resulting in what Horty

2003 calls the *disjunctive approach*. In our example, worlds where I do the taxes and worlds where I do the L&P review are each ranked better than worlds where I do neither and are not ranked either way with respect to each other. So, neither **OUGHT** ϕ nor **OUGHT** ψ is true, but what it is true is **OUGHT** (ϕ or ψ). As Kratzer argues, this is vastly preferable to standard deontic logic, which cannot tolerate incompatible requirements at all. But, Lassiter concludes that is not good enough: “We wanted to find a logic for *ought* and *should* that makes it possible to model moral conflicts, but what we have here is one which simply ignores them.” That’s a bit strong, the classic semantics doesn’t *ignore* conflicts, it adjudicates them. But still, one might reasonably expect that a semantics of *ought* should allow conflicting *oughts* since there is clearly at least a contingent of speakers who allow conflicting *oughts*.

Lassiter (2011) offers his account as one that makes space for conflicting *oughts*. **OUGHT** ϕ is true iff the expected value of ϕ is significantly higher than the average expected value of all the relevant alternatives. This allows two incompatible prejacentes that both ought to be true: as long as there are plenty of alternatives with very low expected value, two alternatives can be significantly higher than the average of all of them.

We do not think this is a good solution. It weakens the semantics of *ought* to an extent that it treats as moral dilemmas cases that are actually far from being dilemmas. I have to apologize to my friend, who lives a few towns away from me. This is a delicate apology, so it is imperative that I look him in the eye while I apologize. Worlds where I apologize in an impersonal way (email, text message, facebook, snail mail, even a phone call without looking him in the eye) are truly beneath contempt. There are two ways I could apologize that satisfy all the relevant requirements: drive over and apologize in person, or make a video call. There are morally irrelevant differences between the two options that might make me prefer one over the other, but as far as what I *ought* to do is concerned, it makes no difference. In this situation, the expected value of ϕ (drive over) and ψ (video call) are both significantly higher than any of the alternatives. Lassiter’s account would now predict that I ought to drive over *and* that I ought to talk to him via video call. And we presume, we should feel there to be a moral dilemma (since that is the intended prediction of his semantics). Neither prediction is correct: there is no moral conflict here and neither of the *ought* sentences is true. Rather what is true is that I ought to take one of the two favored options, and it doesn’t matter which one.

The best we can (expect to) get?

So, it seems like we need to find some middle ground between Lassiter's account which overly cheapens the notion of a dilemma, and the classic semantics, which doesn't really allow dilemmas but always adjudicates them.

There is conventional wisdom about moral dilemmas, usefully explicated by Horty 2003, that we will now quickly review.⁸ There is some unanimity that there is a notion of obligation, traditionally called *prima facie* obligation, that allows for conflicts. In our case, I have to do the taxes and I have to do the L&P review. Each is a genuine *prima facie* obligation of mine. In Kratzer's system, the propositions that I do the tax return and that I do the review will both be in the ordering source. What is controversial is whether after careful consideration of all the *prima facie* obligations an agent is under, the agent's actual/final obligations, often called her *all things considered* obligations, can contain conflicts.

The disjunctive approach, which the classic semantics falls under, adjudicates all *prima facie* conflicts and results in disjunctive all things considered proscriptions, with no conflicts — no conflicts in the sense of there being true all things considered oughts with conflicting prejacent. There may still be consequences of violating a *prima facie* obligation. Even though the classic account predicts that it is not true that I ought to do the taxes and not true that that I ought to do the review, it only being true that I ought to do one or the other, that doesn't mean that there won't be a piper to pay: if I don't do the taxes, I'll be hit with an IRS penalty, if I don't do the L&P review, I'll be in the editor's doghouse.

Before we introduce two ways of modifying the classic semantics so as to make space for genuine conflicting oughts, let us discuss an issue that is repeatedly raised in the literature. It appears that some speakers allow for conflicting oughts but not for conflicting *have tos* or *musts*. Such a judgment is reported, for example, by Harman 1993: p. 185:

[M]ost speakers . . . [find] it quite acceptable to say the following:

- (38) You ought to give C a banana and you ought to give D a banana, but you can't give both of them a banana, so you have to decide.

Most speakers find it less acceptable to say this:

⁸ Meta-ethicists are of course very familiar with the relevant literature. Apart from Horty 2003, we have found the collection of papers in Gowans 1987a useful.

- (39) You have to give C a banana and you have to give D a banana, but you can't give both of them a banana, so you have to decide.

Others with similar judgments are Gowans 1987b: p. 26 and Swanson 2011.

We have to disagree. Moral dilemmas can be talked about with conflicting strong necessity modals as well: "I have a dilemma. My tax return is due midnight tonight. So, I have to work on that. But, my review for L&P is also due tonight. So, I also have to work on *that*. But I *can't* do both. So, I'm screwed, as happens far too often. What should I do?" So, we do not think that the solution will be found in the distinction between weak and strong necessity modals.

There are two ways of modifying the classic semantics so as to make space for conflicting obligation statements. One requires no modification of the semantics proper at all. The idea is to deny that the ordering source in a given context has to be maximal in the sense of collecting all prima facie obligations. Rather, one could say that in the conflict scenarios, there are (at least) three possible salient contextual resolutions of the ordering source parameter. With respect to one ordering source, I have to do the taxes. With respect to another, I have to do the L&P review. And if we put those two ordering sources together into one all-encompassing ordering source, we will make a disjunctive prediction (that I have to do one or the other). With respect to one and the same ordering source, we cannot have a deontic conflict. But if there is more than one salient ordering source, they can pull us into different directions.

The alternative to this "pragmatic" account is to build awareness of a possible multiplicity of ordering sources into the semantics. The idea is then that **OUGHT** ϕ is true with respect to such a multiplicity of ordering sources if (at least) *one* of the ordering sources is such that all the worlds in the modal base that are best by the lights of that ordering source are ϕ -worlds. A special case of such an analysis takes Kratzer's ordering source, which may contain the seeds of conflict, and identifies as relevant (sub-)ordering sources within it, the maximally consistent subsets of propositions. Then **OUGHT** ϕ can be said to be true if ϕ follows from (at least) one of those maximally consistent subsets of propositions from the ordering source. The latter is pretty much the proposal that Horty 2003 makes within a slightly different framework (working with a set of sentences as the prima facie obligations, rather than with propositions in the ordering source as Kratzer does). He

The best we can (expect to) get?

calls this the *conflict account*. The disjunctive account, in contrast, can be formulated as saying that **OUGHT** ϕ is true iff ϕ follows from all maximally consistent subsets of the ordering source.⁹

None of these possible moves have been explored in linguistic semantics, so we're not certain about the prospects of these ways of making space for moral dilemmas. It should be noted, though, that all of these variants of the classic semantics would be monotonic with respect to the prejacent, so they do not depart from the classic picture in that regard. Cariani, whose semantics does not allow moral dilemmas, considers an account that semanticizes reference to a multiplicity of orderings, while maintaining the non-monotonicity he considers a desideratum. So, it is clear that the question of moral dilemmas is independent of the issue of (non-)monotonicity.

As a possible way of deciding which one of these ways to go if one wants to open up the possibility of moral dilemmas within the classic semantics, we could look at the interpretation of negated deontic modals and of permission modals:

- (40) [Context, as before: both the tax return and the L&P are due]
- a. (Luckily,) I don't (really) have to do the taxes [because I (also) have to write the review].
 - b. You are allowed not to do the taxes.

Assuming that the contextual flexibility of resolving which ordering source to be sensitive to extends to negated modals and permission modals, we would predict that these examples have a true reading in the given scenario. It seems to us, however, that these claims are quite clearly false. Under the "semantic" account, (40a) would say that it is not the case that there is at least one (sub-)ordering source which necessitates that I do the taxes. Under that account, negated necessity modals thus have a very strong meaning, explaining the falsity of (40a).

Further, assuming that permission modals are the duals of necessity modals, the semantic account would have them say that there is no (sub-)ordering source which necessitates that the prejacent is false. That again would predict (40b) to be false, which seems right.

⁹ This account is easily transposed into talk using a deontic ordering of worlds. What we need to do is to decompose the ordering underlying the use of a deontic modals into its *maximal subchains*, each of which is a total preorder, rather than the partial preorder with incomparabilities that corresponds to the entire ordering source. In fact, Swanson 2011 proposes just such a semantics for *ought*, which would then allow for conflicting *oughts*.

We can imagine that there are ways of looking at these data differently which would not involve building existential quantification over (sub-)ordering sources into the semantics, but we will not pursue this line of thinking in this paper.

Either way, the conclusion is that moral dilemmas can likely be accommodated within the classic semantics once it is adjusted in some way to allow a context with a multiplicity of (sub-)ordering sources. As far as we can see, the moves here are orthogonal to other features of the classic semantics, in particular its monotonicity.

4 Obligation under uncertainty

The third, and for the purposes of this paper final, challenge to the classic semantics comes from scenarios where there is uncertainty about the facts and likely outcome of possible actions. The problem has been brought to the forefront through the Miners' Puzzle discussed by Kolodny & MacFarlane 2010, but it had been recognized earlier (Lassiter cites Goble 1996, van Rooij 1999 and Levinson 2003, the latter identifying the problem for bouletic attitudes). Here's the case of the miners:

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

On the one hand, it is clear that the best possible outcome is for all the miners to be saved, so we should block the shaft that the miners are in. But we do not have the information about which shaft the miners are in. So, if we randomly block a shaft, we only have a fifty-fifty chance of blocking the right shaft; and there doesn't seem to be any way of improving the odds of blocking the right shaft. So, sensibly, it seems that the following is true:

(41) We ought to block neither shaft. (**UGHT** $\neg(\phi \vee \psi)$)

One aspect of the miners' puzzle is to explain why (41) is considered true. (It has to be said that this is not a part of the puzzle that Kolodny & MacFarlane

The best we can (expect to) get?

(2010) focus on. It is to the credit of Charlow 2010, 2011, Cariani, Kaufmann & Schwager 2011, and Lassiter 2011 that they realize the urgency of this question for our understanding of the semantics of deontic modals.) The other aspect, which we will turn to later, is how deontic modals interact with conditionals in these scenarios of uncertainty.

4.1 What to do when you don't know

Why is this a puzzle? The thought is that the classic semantics fails to correctly predict that (41) is true in the miners' scenario. There are six relevant types of worlds, the cross-product of where the miners are and which, if any, shaft we block.¹⁰ Following Cariani, Kaufmann & Schwager 2011, we can abbreviate those six types of worlds as AA, AB, AN, BA, BB, BN — the first letter standing for which shaft the miners are in, the second letter standing for which shaft (or neither) we block. Clearly, it seems, the two best kinds of worlds are AA and BB: we block the right shaft and save all ten miners. Next best are AN and BN, where we save nine of the ten miners. Worst are AB and BA, worlds where we block the wrong shaft and all miners perish.

(42) AA, BB > AN, BN > AB, BA

So, since the best worlds are AA and BB, it will not be true that we block neither shaft in the best worlds, and thus, (42) is incorrectly predicted to be false.

What is the alternative? Charlow 2010, 2011, Cariani, Kaufmann & Schwager 2011, and Lassiter 2011 present different solutions, which we will briefly discuss in turn. The first two proposals are cast as friendly amendments to the classic semantics, while the last is meant as a wholesale replacement.

Charlow proposes that in this case and specifically for the weak necessity modals (*ought/should*), the ordering is coarsened: "possibilities where we do nothing (and nine are saved) are ranked as highly as those where we block the correct shaft (and all ten are saved)".

¹⁰ Notice that we are assuming here a modal base that does not include the actual world fact of where the miners are. In other words, this is an epistemically flavored modal base. The problem for the classic semantics does not depend on this assumption. Even if we included the miners' actual location (which is unknown to the relevant agents), the classic semantics would still rate the XX world over the XN world — or so it seems.

(43) AA, BB, AN, BN > AB, BA

The idea is that since the only thing we have control over is saving at least nine miners (we cannot ensure that we save 10 miners), worlds where we save nine and worlds where we save ten are given the same rank in the ordering. Charlow thus predicts that it is false that we ought to block either shaft ($\neg\text{OUGHT}(\phi \vee \psi)$). There are at least two worries about his proposal: (i) Charlow's proposal falls short of the main desideratum: the truth of (41).¹¹ (ii) His claim that the issue only arises with weak necessity modals is wrong. In the same scenario, the following are also true: *We must not block either shaft* and *We can't block either shaft*.

Cariani, Kaufmann & Schwager (2011) propose that the modal base of the deontic modal be partitioned into cells according to which action is taken: all the worlds where we block shaft A form one cell, and so do all the B-worlds, and the worlds where we block neither shaft. These three cells are then ranked by the ordering: clearly, the cell where we block neither shaft on average is better than the cells where we block either one of the shafts.

(44) {AN, BN} > {AB, BB}, {AA, BA}

Now, **OUGHT** ϕ is true iff ϕ is true in all of the worlds in the best cell. In our scenario, then, (41) is correctly predicted to be true. Obviously, this proposal constitutes a significant departure from the classic semantics, but as usual, the end justifies the means. Note that the resulting semantics maintains other features of the classic semantics: it is monotonic with respect to the prejacent and it adjusticates deontic dilemmas.

The final non-classic contender is Lassiter's proposal. He follows Goble 1996, van Rooij 1999 and Levinson 2003 in tying deontic modals to the notion of expected value. We are comparing various alternative propositions as to what their expected value is. A reasonable assignment of probabilities and values will rank blocking neither shaft (which assures that nine miners are saved) above blocking either shaft (which has a fifty-fifty chance of killing all miners).¹²

A presupposition in the discussion so far has been that the classic account makes the wrong prediction for (41), which is in turn due to the presuppo-

¹¹ Charlow (2011: Fn.28) is aware of this issue. His fix is criticized by Lassiter (2011: p. 136) as having undesirable consequences

¹² Suspicion: Lassiter wrongly predicts that the following is unambiguously true:

- (i) We ought to block the shaft the miners are in.

The best we can (expect to) get?

sition that the ordering of worlds will have to be as indicated above. This is an incorrect presupposition. The point was already made by Büring 2003 in his commentary on Levinson 2003. It is a simplification to think of the worlds being ranked as being characterized as AA, BB, AN, BN, AB, BA worlds. It is tempting to think that one would necessarily rank AA and BB worlds highest, but the fact is that given the relevant circumstances, which include that we do not know where the miners are, any AA world will be one where we blocked A without knowing whether that was where the miners are and thus running the risk of all killing all of them. It is not obvious that that is a better world than one where nine miners survive because we choose not to run the risk of killing all of them. In fact, in the case at hand, we judge AA worlds to be worse than AN worlds. So, it is entirely plausible that the contextually relevant ordering is as follows:

(45) AN, BN > AA, BB > AB, BA.

With respect to that ordering, the classic semantics makes the correct prediction about (41).

The point can be recast in Kratzer's premise semantics. If the ordering source contained just the propositions {that ten miners are saved, that nine miners are saved}, we would indeed incorrectly rank AA above AN. But as Kratzer lays out in unpublished notes on Cariani, Kaufmann & Schwager 2011, the ordering source might plausibly be something like {"If we know where the miners are, our chosen action yields the optimal outcome for the miners", "If we do not know where the miners are, our chosen action yields a still acceptable outcome for the miners and would not yield a less acceptable outcome if they weren't where they in fact are"}, in which case we obtain the ordering that is appropriate for making the correct prediction for (41). As she writes: "there is nothing in Kratzer's framework that would prevent ordering sources to be information dependent in a way that does justice to the relevant readings of *ought to*. Building information dependence into the ordering source in this way does not give us a rationale for why the assumed norms are what they are. We may still want to invoke decision theory to justify them. But does that mean that the semantics of *ought to* depends on a decision problem parameter? [...] Why pack information about rational decision making into the meaning of modals?" This echoes

The proposition " λw . we block in w the shaft the miners are in in w " has a higher expected value than any alternative (in particular, blocking neither shaft or blocking the wrong shaft). Nevertheless, we believe that (i) has a reading where it is false.

Büring’s complaint about Levinson’s argumentation: “I contend that there is no need to build the rational behavior of determining what is desirable into the semantics of verbs of desire like *want*.”

The point is that it is perfectly compatible with the classic semantics to have the ordering of worlds be sensitive to the information available in each of the worlds ordered. We don’t know where the miners are. If we take this to be one of the defining circumstances of the modal base, then each of the worlds in the modal base will be one where we don’t know where the miners are. If we then work with an ordering that favors worlds where we don’t do anything that as far as we know in that world, might kill more miners than it saves, we obtain the kind of ordering that will make it true that we ought block neither shaft.

If we modulate the circumstances just a little bit, we obtain different predictions. Imagine that we have an hour to come up with our course of action, and imagine that it is at least possible that we can find out in that time where the miners are. Then, it becomes true that we ought to find out where the miners are and then block the shaft that they’re in.

None of this contradicts the classic semantics.¹³

4.2 Iffy oughts

It might be complained that this account does not cover the facts about conditional oughts in the Miners case. As Kolodny & MacFarlane 2010 point out, the following two conditionals are instantly judged as true in the scenario:

- (46) If the miners are in Shaft A, we ought to block Shaft A.
 (47) If the miners are in Shaft B, we ought to block Shaft B.

The classic semantics, as filled out just now, together with Kratzer’s Thesis about conditionals (that is, the thesis that *if*-clauses “restrict” operators,

¹³ One thing we will not explore in this paper is the exact notion of information or knowledge sensitivity that enters into subjective deontic modals. It is likely that in these uses deontic modals will display behavior similar to epistemic modals: sometimes more than just the actual knowledge of the relevant agents is crucial, what matters might be the information available to them in some sense, or the information in their “epistemic reach”, to borrow Andy Egan’s phrase (Egan 2007). So, for example, it seems that what the rescuers ought to do depends not just on what they actually know but on information that is in their reach, such as the easily accessible logbook which should reveal what shaft the miners are in: in that case, they ought to block the right shaft even if they culpably remain ignorant of what’s in the logbook. Thanks to Angelika Kratzer for discussion on this point.

The best we can (expect to) get?

which in the case of modals means that they narrow the modal base to those worlds verifying the antecedent) predicts that (46) and (47) should be just as false as the unconditional “We ought to block Shaft A” and “We ought to block Shaft B”. The reason is that no matter where the miners are, it is given in the scenario that we don’t know where they are and the ordering source, as just explained, is insensitive to where the miners are. Even if the miners are in Shaft A, we don’t know where they are and so blocking Shaft A is risky relative to our information. So, how come (46) and (47) are judged as true?

It is tempting to argue that these conditionals involve an “objective” *ought*. Part of what underwrites (41) is the fact that we don’t know where the miners are. And if that is given, then the mere assumption that they are in Shaft A makes no difference: we shouldn’t block either shaft. The following examples make the ignorance assumption explicit:

- (48) Given that we don’t know where the miners are, we ought to block neither shaft.
- (49) Given that we don’t know where the miners are, if the miners are in Shaft A, we ought to block Shaft A.
- (50) Given that we don’t know where the miners are, if the miners are in Shaft B, we ought to block Shaft B.

It seems clear that (48) remains as true as (41) and really is just a more explicit way of saying the same thing. But there is something deeply problematic about (49) and (50). Once it is given that we don’t know where the miners are, it doesn’t matter where they are for what we ought to do in the sense of maximizing our chances of doing the right thing. In fact, the correct thing to say while maintaining the same resolution of *ought* is what Cariani, Kaufmann & Schwager 2011 mention in a footnote:

- (51) If the miners are in shaft A, we (still) ought to block neither shaft, for their being in shaft A doesn’t mean that we know where they are. Indeed, no matter where the miners are, we ought to block neither shaft.

Therefore, it seems clear that Kolodny & MacFarlane’s conditionals involve a different resolution of the modal parameters. In particular, we might think that they involve an “objective” kind of obligation where we do not take into consideration what we know but just what the facts are. Under that resolution, the facts about the actual location of the miners will make a

difference to what ought to be done. It would also be easy to explain why this resolution of the deontic modal is the natural one for the conditionals: against the background assumption that we don't know where the miners are, restricting the modal base via an *if*-clause will have no effect on the truth of the modal; this is made explicit in (51). This would appear to be a fairly useless contribution to the conversation. Under the "objective" resolution, however, the *if*-clauses do make a difference and so that resolution is favored.

This may not be the whole story, though. Several authors have suggested that what goes on in these conditionals is that the *if*-clause modifies the relevant information state that affects the information-sensitive ordering.¹⁴ This would be non-classic and if right, be fatal to the classic semantics. While we think that the re-analysis of the Kolodny & MacFarlane conditionals we just sketched makes the right predictions for those specific cases, the general picture is more complicated. To see this, we need to add an additional wrinkle to the miners scenario.

Imagine there are three shafts: A, B, and C. We don't know where the miners are. If we block the right shaft, all miners are safe. If we do nothing, two miners die. We can blow up Shaft A, which would of course kill all miners if they are in A, but if they're not, then blowing up Shaft A and not blocking either B or C will mean that only one miner dies. So, in our maximally ignorant information state, we ought to block none of the shafts. In an objective sense, we should block the shaft the miners are in. Now, consider the following conditional:

(52) If they are not in A, we ought to blow A up.

I can hear this as true. Our analysis so far does not predict this true reading. So, this might be thought to be a knockdown argument for the non-classic analysis of *iffy* oughts (where they modify the non-objective information state that **UGHT** is sensitive to).

Here's what I think. (52) is shorthand, enthymematic for (53).

(53) If we learn that they are not in A, we ought to blow A up.

That is, I think that the conditional does take us to an information state where it is known that they are not in A. But it doesn't do so via a semantic mechanism that affects a contextual information state parameter. Rather, the conditional is understood as not just restricting the modal base to worlds

¹⁴ In addition to the works already cited, see Silk 2011.

The best we can (expect to) get?

where the antecedent is true but as moving us to worlds where we have learned that the antecedent is true.

This is not the standard way of construing conditionals. As pointed out long ago by van Fraassen (1980: p. 503), attributing the observation to Thomason, if conditionals always worked that way, we wouldn't be able to make sense of examples like

(54) If my partner is cheating on me, I'll never find out.

Having said that, it seems that such a construal of conditionals is actually possible:

(55) If my partner is cheating on me, I'll be surprised.

(55) only makes sense if we're being transported to scenarios in which not only is my partner cheating on me but also I learn that this is so. Otherwise, of course, I couldn't be surprised.

So, my claim is that as much as (52) is taken to be true, this is because it is interpreted in this enthymematic way. Here are a couple of considerations in favor of this view. First, the following variants on the original Kolodny & MacFarlane conditionals seem far less acceptable (thanks to Angelika Kratzer for devising these examples):

- (56) a. If the miners are in shaft A, we ought to get sandbags right away and block it.
b. If the miners are in shaft A, we ought to act fast and block it before the miners suffocate.
c. If the miners are in shaft A, let's get sandbags and block it!

If we are right and these examples are less acceptable as true, then we have an indication that in the original examples we were actually dealing with the assumption that we will learn that the antecedent is true. The examples in (56) are constructed to be about what is right now the case and what we ought to do this very instant and thus make it harder to let the conditional take us to a future situation in which we have learned that the antecedent is true.

Second, let me point out a crucial feature of all the examples we have discussed in this section: they are all formulated in the first person plural. I suspect that this smuggles in a conflation between the information state that the deontic modal is sensitive to and the information state of the speaker(s).

Let's try a scenario that doesn't conflate the roles.

Imagine that we are observers of the would-be rescuers of the stranded miners and charged with evaluating their performance (perhaps, to decide whether they should be hired as a rescue task force for future mining disasters). Suppose we have no way of communicating with the rescuers. We can now truthfully say the following:

(57) They ought to block neither shaft.

Note that (57) seems true whether or not *we*, the observers, know where the miners are. Blocking neither shaft is the right thing to do. If the rescuers chose to block the correct shaft, we might feel that it's a good thing that they did that, because it saved all the miners. But we would not recommend them as a rescue team for future disasters, since they ran a huge risk.

Now, consider the following conditional in this scenario:

(58) If the miners are in Shaft A, the rescuers ought to block A.

Two observations: (i) if we know where the miners are, (58) is pretty weird, presumably because of the iffiness signaled: the conditional conveys that the speaker is open to the antecedent being true or false; (ii) if we don't know where the miners are, (58) is felicitous but it is pretty clearly false: whether or not the miners are in Shaft A, whether or not *we*, the observers, learn that the miners are in Shaft A, has no effect on what is the right thing to do on the part of the rescuers. Since they don't know where the miners are, they ought to block neither shaft. Accounts that predict that the *if*-clause updates the information state that **OUGHT** is sensitive to will wrongly predict that (58) should be true in the given scenario.

In sum, the interaction of **OUGHT** with conditionals, even in scenarios of uncertainty, does not force us to abandon classic semantics, and if the last two arguments are correct, it would in fact be wrong to abandon classic semantics.

5 Conclusion

For each of the challenges we have discussed, we have argued that the classic semantics, sometimes elaborated or modified, can meet those challenges. As we said in the beginning, this is all very programmatic, but the point is that the classic semantics is not dead yet and in fact, may be the best we can (expect to) get.

The best we can (expect to) get?

References

- Aloni, Maria. 2007. Free choice, modals, and imperatives. *Natural Language Semantics* 15(1). 65-94. <http://dx.doi.org/10.1007/s11050-007-9010-2>.
- Alonso-Ovalle, Luis. 2006. *Disjunction in alternative semantics*. University of Massachusetts at Amherst PhD thesis. http://alonso-ovalle.net/index.php?page_id=28.
- Asher, Nicholas. 1987. A typology for attitude verbs and their anaphoric properties. *Linguistics and Philosophy* 10(2). 125-197. <http://dx.doi.org/10.1007/BF00584317>.
- Büring, Daniel. 2003. To want it is to want to be there: a note on Levinson (2003). ms, UCLA, paper for workshop on *Division of Linguistic Labor* at the Château de la Bretesche near Nantes, France. <http://www.linguistics.ucla.edu/general/Conf/LaBretesche/papers/buring.pdf>.
- Cariani, Fabrizio. 2011. *Ought* and resolution semantics. ms, to appear in *Noûs*. <http://cariani.org>.
- Cariani, Fabrizio, Stefan Kaufmann & Magdalena Schwager. 2011. Deliberative modality under epistemic uncertainty. ms. http://semanticsarchive.net/Archive/TM5ZmM2Z/epi_delib_final.pdf.
- Charlow, Nate. 2010. Restricting and embedding imperatives. In Maria Aloni, Harald Bastiaanse, Tiki de Jager & Katrin Schulz (eds.), *Logic, language and meaning: 17th Amsterdam Colloquium, Amsterdam, The Netherlands, December 16-18, 2009, revised selected papers* (Lecture Notes in Artificial Intelligence 6042), 223-233. Springer. http://dx.doi.org/10.1007/978-3-642-14287-1_23.
- Charlow, Nate. 2011. What we know and what to do. *Synthese*. <http://dx.doi.org/10.1007/s11229-011-9974-9>.
- Chierchia, Gennaro, Danny Fox & Benjamin Spector. 2012. The grammatical view of scalar implicatures and the relationship between semantics and pragmatics. To appear in von Stechow, Maienborn & Portner 2012, Chapter 87. http://semanticsarchive.net/Archive/WMzY2ZmY/CFS_EmbeddedSIs.pdf.
- Crnič, Luka. 2011. *Getting even*. Cambridge, MA: Massachusetts Institute of Technology PhD thesis. <http://pluto.huji.ac.il/~crnic/crnic-diss-11.pdf>.
- Danielsson, Sven. 1968. *Preference and obligation*. University of Uppsala, Sweden: Filosofiska föreningen, Department of Philosophy Series No. 7.
- Egan, Andy. 2007. Epistemic modals, relativism, and assertion. *Philosophical Studies* 133(1). 1-22. <http://dx.doi.org/10.1007/s11098-006-9003-x>.

- von Fintel, Kai. 1999. NPI licensing, Strawson entailment, and context dependency. *Journal of Semantics* 16(2). 97-148. <http://dx.doi.org/10.1093/jos/16.2.97>.
- von Fintel, Kai. 2001. Counterfactuals in a dynamic context. In Michael Kenstowicz (ed.), *Ken Hale: a life in language*, 123-152. MIT Press. <http://mit.edu/fintel/fintel-2001-counterfactuals.pdf>.
- von Fintel, Kai & Sabine Iatridou. 2008. How to say *ought* in Foreign: the composition of weak necessity modals. In Jacqueline Guéron & Jacqueline Lecarme (eds.), *Time and modality* (Studies in Natural Language and Linguistic Theory 75), 115-141. Springer. <http://dx.doi.org/10.1007/978-1-4020-8354-9>.
- Føllesdal, Dagfinn & Risto Hilpinen. 1971. Deontic logic: an introduction. In *Deontic logic: Introductory and systematic readings*. Dordrecht: Reidel.
- van Fraassen, Bas C. 1980. Critical notice of Brian Ellis, *Rational Belief Systems*. *Canadian Journal of Philosophy* 10(3). 497-511.
- van Fraassen, Bas C. 1972. The logic of conditional obligation. *Journal of Philosophical Logic* 1(3-4). 417-438. <http://dx.doi.org/10.1007/BF00255570>.
- Frank, Anette. 1996. *Context dependence in modal constructions*. Universität Stuttgart PhD thesis. <http://www.cl.uni-heidelberg.de/~frank/papers/header.pdf>.
- Geurts, Bart. 2004. On an ambiguity in quantified conditionals. ms, University of Nijmegen. <http://ncs.ruhosting.nl/bart/papers/conditionals.pdf>.
- Gillies, Anthony S. 2007. Counterfactual scorekeeping. *Linguistics and Philosophy* 30(3). 329-360. <http://dx.doi.org/10.1007/s10988-007-9018-6>.
- Goble, Lou. 1996. Utilitarian deontic logic. *Philosophical Studies* 82(3). 317-357. <http://dx.doi.org/10.1007/BF00355312>.
- Gowans, Christopher W. (ed.). 1987a. *Moral dilemmas*. New York & Oxford: Oxford University Press.
- Gowans, Christopher W. 1987b. The debate on moral dilemmas. In Christopher W. Gowans (ed.), *Moral dilemmas*, 3-33. New York & Oxford: Oxford University Press.
- Hansson, Bengt. 1969. An analysis of some deontic logics. *Noûs* 3(4). 373-398. <http://dx.doi.org/10.2307/2214372>.
- Hare, R.M. 1967. Some alleged differences between imperatives and indicatives. *Mind* 76(303). 309-326. <http://dx.doi.org/10.1093/mind/LXXVI.303.309>. <http://www.jstor.org/pss/2251842>.

The best we can (expect to) get?

- Harman, Gilbert. 1993. Stringency of rights and *ought*. *Philosophy and Phenomenological Research* 53(1). 181–185. <http://dx.doi.org/10.2307/2108064>.
- von Heusinger, Klaus, Claudia Maienborn & Paul Portner (eds.). 2012. *Semantics: an international handbook of natural language meaning*, vol. 3. Berlin: de Gruyter.
- Horty, John F. 2003. Reasoning with moral conflicts. *Noûs* 37(4). 557–605. <http://dx.doi.org/10.1046/j.1468-0068.2003.00452.x>.
- Iatridou, Sabine & Hedde Zeijlstra. 2011. Negation, polarity and deontic modals. ms, MIT and University of Amsterdam. <http://ling.auf.net/lingBuzz/001431>.
- Jackson, Frank & Robert Pargetter. 1986. Oughts, options, and actualism. *The Philosophical Review* 95(2). 233–255. <http://dx.doi.org/10.2307/2185591>.
- Kamp, Hans. 1973. Free choice permission. *Proceedings of the Aristotelian Society, New Series* 74. 57–74. <http://www.jstor.org/stable/4544849>.
- Kanazawa, Makoto. 1991. Puzzles about conjunction and disjunction in certain intensional contexts. In Akira Ikeya (ed.), *The sixth Japanese-Korean joint conference on formal linguistics*, 113–132. Logico-Linguistic Society of Japan. <http://dx.doi.org/2065/11780>.
- Klecha, Peter. 2011. Positive and conditional semantics for gradable modals. Talk given at *Sinn und Bedeutung* 16.
- Kolodny, Niko & John MacFarlane. 2010. Ifs and oughts. *The Journal of Philosophy* 108(3). 115–143.
- Kratzer, Angelika. 1977. What *must* and *can* must and can mean. *Linguistics and Philosophy* 1(3). 337–355. <http://dx.doi.org/10.1007/BF00353453>.
- Kratzer, Angelika. 1981. The notional category of modality. In Hans-Jürgen Eikmeyer & Hannes Rieser (eds.), *Words, worlds, and contexts: new approaches in word semantics* (Research in Text Theory 6), 38–74. Berlin: de Gruyter.
- Kratzer, Angelika. 1991. Modality. In Arnim von Stechow & Dieter Wunderlich (eds.), *Semantics: an international handbook of contemporary research*, 639–650. Berlin: de Gruyter. <https://udrive.oit.umass.edu/kratzer/kratzer-modality.pdf>.
- Kratzer, Angelika. 2010. Papers on modals and conditionals. unpublished book ms.
- Kratzer, Angelika & Junko Shimoyama. 2002. Indeterminate pronouns: the view from Japanese. Paper presented at the 3rd Tokyo Conference on

- Psycholinguistics March 2002. <http://semanticsarchive.net/Archive/WEwNjc4Z/>.
- van Kuppevelt, Jan. 1996. Inferring from topics: Scalar implicatures as topic-dependent inferences. *Linguistics and Philosophy* 19(4). 393-443. <http://dx.doi.org/10.1007/BF00630897>.
- von Kutschera, Franz. 1974. Normative Präferenzen und bedingte Gebote. In Hans Lenk & Jörg Berkemann (eds.), *Normenlogik: Grundprobleme der deontischen Logik* (UTB 414), 137-165. Pullach bei München: Verlag Dokumentation. <http://epub.uni-regensburg.de/12608/>.
- Lassiter, Daniel. 2011. *Measurement and modality: the scalar basis of modal semantics*. New York University PhD thesis. <http://www.stanford.edu/~danlass/Lassiter-diss-Measurement-Modality.pdf>.
- Levinson, Dmitry. 2003. Probabilistic model-theoretic semantics for *want*. *Semantics and Linguistic Theory (SALT)* 13. 222-239. <http://elanguage.net/journals/index.php/salt/article/viewFile/13.222/1823>.
- Lewis, David. 1973. *Counterfactuals*. Oxford: Blackwell.
- Lewis, David. 1974. Semantic analyses for dyadic deontic logic. In Sören Stenlund (ed.), *Logical theory and semantic analysis: Essays dedicated to Stig Kanger on his fiftieth birthday*, 1-14. Dordrecht: Reidel.
- Moss, Sarah. 2010. On the pragmatics of counterfactuals. *Noûs*. <http://dx.doi.org/10.1111/j.1468-0068.2010.00798.x>. Early View (Online Version of Record published before inclusion in an issue).
- Nauze, Fabrice. 2008. *Modality in typological perspective*. Universiteit van Amsterdam PhD thesis. <http://www.illc.uva.nl/Publications/Dissertations/DS-2008-08.text.pdf>.
- Ninan, Dilip. 2005. Two puzzles about deontic necessity. In Jon Gajewski, Valentine Hacquard, Bernard Nickel & Seth Yalcin (eds.), *New work on modality* (MIT Working Papers in Linguistics 51), 149-178. Department of Linguistics & Philosophy, MIT. <http://semanticsarchive.net/Archive/WZINTgoY/NinanDeontic.pdf>.
- van Rooij, Robert. 1999. Some analyses of pro-attitudes. In Henriëtte de Swart (ed.), *Logic, game theory and social choice*, 534-548. Tilburg: Tilburg University Press. <http://staff.science.uva.nl/~vanrooij/Desire.pdf>.
- Ross, Alf. 1941. Imperatives and logic. *Theoria* 7(1). 53-71. <http://dx.doi.org/10.1111/j.1755-2567.1941.tb00034.x>. Reprinted with only minor editorial changes as Ross 1944.
- Ross, Alf. 1944. Imperatives and logic. *Philosophy of Science* 11(1). 30-46. <http://dx.doi.org/10.1086/286823>.

The best we can (expect to) get?

- Sauerland, Uli. 2004. On embedded implicatures. *Journal of Cognitive Science* 5. 107-137. <http://homepage.mac.com/yasa/uli/journals/JCSimplicatures.pdf>.
- Silk, Alex. 2011. Evidence sensitivity in deontic modals. ms, University of Michigan. http://www-personal.umich.edu/~asilk/Alex_Silk/home_files/silk%20evidence-sensitivity.pdf.
- Simons, Mandy. 2005. Dividing things up: the semantics of *or* and the modal/*or* interaction. *Natural Language Semantics* 13(3). 271-316. <http://dx.doi.org/10.1007/s11050-004-2900-7>.
- Swanson, Eric. 2011. On the treatment of incomparability in ordering semantics and premise semantics. *Journal of Philosophical Logic* 40(6). 693-713. <http://dx.doi.org/10.1007/s10992-010-9157-z>.
- Wedgwood, Ralph. 2006. The meaning of *Ought*. In Russ Shafer-Landau (ed.), *Oxford studies in metaethics*, vol. 1, 127-160. Oxford University Press. <http://users.ox.ac.uk/~mert1230/meaningofought.ltr.pdf>.
- Zvolenszky, Zsófia. 2002. Is a possible-worlds semantics of modality possible? A problem for Kratzer's semantics. *Semantics and Linguistic Theory (SALT)* 12. <http://web.gc.cuny.edu/dept/lingu/liba/papers/Zsolenszky2002.pdf>.