SUBJECTIVE *Ought*

JENNIFER CARR

University of Illinois at Urbana-Champaign

THIS paper concerns the subjective *ought* in natural language. If you think the apple is poisoned, there's a sense in which you ought not eat the apple—even if, unbeknownst to you, the apple isn't poisoned. This *ought*, the subjective *ought*, isn't just sensitive to sources of value in the world: it's also sensitive to what information is available. The objective *ought*, by contrast, is insensitive to knowledge and ignorance: it's the *ought* from a God's-eye view.

The subjective and objective *ought* and related expressions (*should*, arguably *must* and *have to*, attitude verbs like *want* and *need*, comparatives like *better*, and so on) aren't technical terms in philosophy. They are a part of natural language, the language we use to express intuitions in ethics, philosophy of action, and decision theory. When we investigate the relations between what we ought to do and considerations like what morality or prudence requires, what we want or intend, what we're able to do, and so on, our claims are expressed in natural language. They reflect the logical structure of natural language. And our theorizing often tacitly makes assumptions about the entailment relations between various considerations and what we ought to do.

So we should get clear on how *ought* and related expressions work in natural language. The subjective *ought*, in particular, is not well understood. The recent literature on the Miners Puzzle, initiated by Kolodny and MacFarlane (2010) and developed by Charlow (2013), Cariani, Kaufmann, and Kaufmann (2013), Dowell (2012), von Fintel (2012), and Silk (2014a), discredits a widespread misconception about the subjective *ought*: that *ought* claims in natural language must obey classical inference rules like *modus tollens*. In fact, subjective *ought* provides a wealth of examples that show that those rules aren't valid. In Section 1, I summarize Kolodny and MacFarlane's argument that indicative conditionals don't obey *modus ponens*, and I provide some new empirical and theoretical arguments against the validity of classical inference rules for natural language conditionals.

Contact: Jennifer Carr <jrcarr@illinois.edu>

Kolodny and MacFarlane's example also casts doubt on the orthodox theory of modals and conditionals, Kratzer's semantics (1981; 1991), even though Kratzer's semantics does not validate classical inference rules. Charlow (2013) and Cariani et al. (2013) suggest conservative amendments to the orthodox view that are meant to accommodate Miners Puzzle phenomena. In Section 2, I argue that these views also face counterexamples.

I argue in Section 4 that these counterexamples are manifestations of a deeper problem for the orthodox semantics and its recent variants: they all inadvertently build substantive (and, moreover, unattractive) normative assumptions into the semantics of modals. These normative assumptions come in the form of (structural) decision rules: they tell us how to go from some objective body of values to a verdict about what subjectively we ought to do, given our limited information. The fact that these decision rules are unattractive explains why many of the resulting predictions are intuitively wrong. It's only once we've seen how the orthodox picture and its refinements build in these normative commitments that the motivations for my own view become clear. My central claim, in Section 5, is that the subjective *ought* is sensitive not only to priorities and to information, but also to norms for the pursuit of priorities given limited information (in other words, decision rules). Instead of building these norms into the semantics, I argue, we should let them be determined by context. Finally, some of these norms require rethinking how we represent the effects of information and values on the subjective *ought*.

1. Subjective *Ought* Doesn't Obey Classical Inference Rules

This paper focuses on the now well-known Miners Puzzle. The puzzle illustrates two separate and important points. The first is that the subjective *ought* can give rise to counterexamples to classical inference rules. The second is that the example exposes a problem for the orthodox semantics for subjective *ought*. I address the first point in this section. The second is the focus of the rest of the paper.

1.1. The Miners Puzzle

A first observation: claims involving subjective *ought* can generate counterexamples to classical inference rules. Kolodny and MacFarlane (2010) introduce the following counterexample to proof by cases:¹

1. Kolodny and MacFarlane argue that the Miners Puzzle is a counterexample to modus ponens for indicative conditionals. Their argument presupposes a specific, static interpretation of validity. Let $\llbracket \cdot \rrbracket^c$ denote a valuation function from sentences uttered in any context c to possible worlds propositions. Then:

The Miners Puzzle. “Ten miners are trapped either in shaft A or in shaft B, but we do not know which. [They’re 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.”²

In this context, the sentences in (1) seem to be true:

- (1)
 - 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.

The puzzle is that according to classical inference rules, (1b), (1c), and (1d) entail (2):

- (2) Either we ought to block shaft A or we ought to block shaft B.

On its face, this example shows that classic inference rules like modus ponens or proof by cases are not valid in natural language.³

1.2. Objection #1: Equivocation

It’s sometimes thought that these examples can be explained away by appeal to *ambiguity*: the uses of *ought* at work exhibit the subjective/objective *ought* ambiguity. While (1a) contains the subjective *ought*, (1b) and (1c) contain the objective *ought*. So, the objection goes, their apparent consistency relies on equivocation. It doesn’t prove that the relevant classical inference rules aren’t valid.⁴

This objection isn’t successful: the modus ponens violation doesn’t rely on equivocation. My argument will look roughly like this: even if (1b) and (1c) weren’t true on the subjective reading, the puzzle reappears in cases that use the

VALIDITY: $\phi_1, \dots, \phi_n \models \psi$ iff $(\llbracket \phi_1 \rrbracket^c \cap \dots \cap \llbracket \phi_n \rrbracket^c) \subseteq \llbracket \psi \rrbracket^c$

Other, dynamic notions of validity, e.g., in Yalcin (2012), suggest the problem afflicts proof by cases and modus tollens but not modus ponens. I prefer to remain noncommittal about which notion of validity is relevant and so will argue against “classical decision rules.” (N.B. I use the bare plural “classical inference rules” existentially, not generically.) See Bledin (2014) for discussion of how these linguistic phenomena should affect our views on validity.

2. Kolodny and MacFarlane (2010: 115), who take this example from Parfit (1988), who attributes it to Regan (1980).

3. Note: none of this is meant to show that modus ponens or any of the other rules is invalid for *material conditionals*. Rather, this is further evidence that indicatives are not material conditionals.

4. Dowell (2012) defends this sort of solution to the Miners Puzzle.

subjective *ought* throughout, and in cases where instead of *ought* we use other, unambiguous information-sensitive expressions. So the ambiguity explanation doesn't solve the general problem, and a solution to the general problem will make appeal to ambiguity otiose.

We can generate Miners Puzzle-like cases where the subjective reading of *ought* is both clearly available and truth-conditionally distinct from the objective reading. Because in these cases the *ought* is subjective throughout, the objective/subjective *ought* distinction cannot save classical inference rules. An example:

The New Miners Puzzle. We're in the same situation as before. Furthermore: it's either Monday or Tuesday but we have no idea which. If it's Monday, then it's 95% likely that the miners are in shaft A. If it's Tuesday, then it's 95% likely that they're in shaft B.

- (3) a. It's either Monday or Tuesday.
- b. If it's Monday, we ought to block shaft A.
- c. If it's Tuesday, we ought to block shaft B.
- d. We ought to block neither shaft.

Here the objective and subjective interpretations have different truth conditions. In the New Miners Puzzle scenario, the speaker is in a position to know that (3b) and (3c) are true. She wouldn't be in a position to know that fact if the *ought* were objective, because she cannot ignore the possibility that she's found herself in the 5% likely case, where the miners are in the opposite shaft. So (3b) and (3c) must both use the subjective *ought*.

So appeal to this sort of ambiguity will not protect modus ponens or other classic inference rules from counterexamples.

It's worth noting that, in fact, we have no reason to think that the original Miners Puzzle conditionals (1b) and (1c) use the objective *ought*. The conditionals in the New Miners Puzzle, (3b) and (3c), show something important about the behavior of subjective *ought* when it appears in the consequent of a conditional. The subjective reading of *ought* is not merely sensitive to the information in the agent's (or anyone else's) *actual* epistemic state. Rather, it's sensitive to that information supplemented with the information from the antecedent of the conditional. In the next subsection, I'll show that this is a general feature of how information-sensitive operators interact with indicative conditionals.⁵

Now let's reconsider (1b): *If the miners are in shaft A, we ought to block shaft A.* If we evaluate the subjective *ought* claim relative to the speaker's information *plus*

5. Those familiar with Kratzer's semantics for modals and conditionals may believe that her account satisfactorily captures this fact. Kolodny and MacFarlane argue that, in the case of deontic modals, it does not; see Section 2.

the information in the antecedent, then it comes out true. After all, that body of information “knows” where the miners are. So, perhaps surprisingly, there’s no good reason to posit equivocation in the original case. On the subjective *ought* reading, the original Miners’ Puzzle conditionals are perfectly true. This is just one of those places where the subjective and objective *ought* happen to coincide.

1.3. Objection #2: Error Theory

A naive objection to the New Miners Puzzle: the conditionals (3b) and (3c) must be false. Even if it is Monday, given the credences that the speaker has in the context of utterance, she should block neither shaft. Philosophers of language and others familiar with the terrain of semantic theories of conditionals may find this objection obviously misguided. But explaining the way in which it is misguided will help to clarify why, in general, the Miners Puzzle causes problems even for semantically sophisticated theories, including the semantic orthodoxy.

The so-called “Ramsey test” is a good heuristic for predicting the acceptability of indicatives (at least on some interpretations). The Ramsey test (1931) involves temporarily imagining the information in the antecedent of an indicative conditional added to one’s body of knowledge and then, under this supposition, evaluating the consequent.

Many who reject these conditionals seem to think that if they have true subjective readings, it must be possible to gloss them as saying:

- (4) If it’s Monday, then *given our current evidence*, we ought to block shaft A.
- (5) If it’s Tuesday, then *given our current evidence*, we ought to block shaft B.

But that’s not how embeddings in natural language work. While the subjective *ought* is by definition information-sensitive, the information to which it is sensitive need not be precisely the information the speaker (or anyone else) actually has. Indicatives’ antecedents can have an effect on what information is relevant. And so the glosses with the additional clause *given our current knowledge* cannot capture the meanings of the Miners Puzzle and New Miners Puzzle conditionals.

Consider another information-sensitive expression, *probably*. When the operator appears in the consequent of the conditional, the information state that affects the modal isn’t the speaker’s epistemic state. It isn’t *anyone’s* epistemic state—at least not by default. It’s a *modification* of the speaker’s epistemic state (or the conversational common ground) so that it includes the information in the conditional’s antecedent. There’s no reason to expect those operators’ deontic cousin, subjective deontic *ought*, to behave any differently—especially in light of clear empirical evidence that it behaves in just the same way.

Here's one piece of evidence: sentences involving *probably* and epistemic *must* also give rise to violations of classical inference rules.⁶

Suppose that it's either Monday or Tuesday; we have no idea which. And suppose the miners could be in shaft A, shaft B, or neither shaft: the likelihood of each is $\frac{1}{3}$. They're in shaft A only if it's Monday and in shaft B only if it's Tuesday. If it's Monday, it's $\frac{2}{3}$ likely that the miners are in shaft A and if it's Tuesday, it's $\frac{2}{3}$ likely that they're in shaft B.

- (6)
- It's either Monday or Tuesday.
 - If it's Monday, they're probably in shaft A.
 - If it's Tuesday, they're probably in shaft B.
 - They're probably not in shaft A.
 - They're probably not in shaft B.

Since the subjective *ought* produces the same pattern of counterexamples as other information-sensitive expressions, we should expect a common explanation. And that explanation will not be able to appeal to a subjective/objective equivocation.

The observation that information-sensitive operators are sensitive to bodies of abstract information that are manipulated under embeddings—for example, restricted by conditional antecedents—is not a novel observation. It's necessary to make sense of ordinary sentences like *if it's raining, the streets must be wet* (in circumstances where the streets might not be wet). I emphasize this point for two reasons: first, it is not obvious to nonspecialists. Second, as we'll see in Section 2, it creates a general theoretical challenge to Kratzer's semantics. While Kratzer's semantics by design accommodates information sensitivity in the background assumptions governing the interpretation of deontic modals, it cannot accommodate information-sensitivity in fixed priorities.

1.4. Objection #3: Ellipsis

An idea that is commonly floated is that we can explain away the pattern of behavior I've described for conditionals with subjective *ought* and other information-

6. Yalcin (2012) independently reached this conclusion about *probably* on similar grounds.

sensitive expressions by denying the literal and non-elliptical truth of (1b) and (1c). Kratzer, in unpublished notes,⁷ and von Fintel (2012) both defend the following hypothesis:⁸ the Miners Puzzle conditionals elliptically express:

- (7) a. If the miners are in shaft A *and we {know/learn} that they are*,
 then we should block shaft A.
 b. If the miners are in shaft B *and we {know/learn} that they are*,
 then we should block shaft B.

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

There are several serious problems with this account.⁹ Here I will only develop one. First, as Thomason observed,¹⁰ conditionals don't generally allow for this sort of addition of *and we {know/learn} it* to the antecedent. Compare:

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

This leads to a fundamental problem with this defense: we can generate cases just like the Miners Puzzle case where it's clear that the conditionals *if ϕ , ψ* cannot be glossed as *if ϕ and we learn ϕ , ψ* . There are no easy solutions available to explain away these cases.

Suppose I'm trying not to find out what my mother got me for my birthday. If an action will probably ($> 50\%$ likely) lead to me knowing what she got me, I don't want to do it. (Of course, if I already know what she got me, this consideration will be moot.) I also want to find my keys, which are either in the closet or in the drawer—but only if doing so doesn't conflict with the first desire. She either got me a necklace or a sweater with equal likelihood. If she got me a necklace, it's $\frac{2}{3}$ likely to be in the drawer, and if she got me a sweater, it's $\frac{2}{3}$ likely that it's in the closet. So there is only $\frac{1}{3}$ probability that there's a present in the drawer and only $\frac{1}{3}$ probability that there's a present in the closet.

7. These notes are summarized in (von Fintel 2012).

8. Note that Kratzer and von Fintel do not intend to defend classical inference rules, which Kratzer's semantics does not validate. They instead propose this hypothesis in defense of the Kratzer's semantics, which Kolodny and MacFarlane's example also afflicts; see Section 2.

9. One objection, which I won't elaborate on but which Fabrizio Cariani (personal communication) has also noted, is that in some Miners Puzzle contexts the conditionals suffer presupposition failure. In some contexts we *know* we don't know and won't learn where the miners are.

10. Attributed to Thomason in van Fraassen (1980).

ND: necklace in drawer

$\overline{ND}$: necklace not in drawer

SC: sweater in closet

$\overline{SC}$: sweater not in closet

Take a desire-based, instrumental reading of *should*. The following sentences are true:

- (9)
- Either she got me a sweater or she got me a necklace.
 - If she got me a sweater, I shouldn't look in the closet.
 - If she got me a necklace, I shouldn't look in the drawer.
 - It's not the case that I shouldn't look in the drawer and it's not the case that I shouldn't look in the closet.¹¹
 - I should look in the drawer or in the closet.

All of the *shoulds* in (9) can easily be interpreted subjectively. The priorities at stake don't change between the sentences. In all of the sentences, *should* will naturally access a modal background where I don't know what my mother got me. Glosses that insert *and I {know/learn} of it* into the antecedents of the conditionals are clearly truth conditionally different from the original conditionals—and in this case, they are clearly false.

So the elliptical-*know/learn* account cannot explain away these sorts of cases. The phenomenon is simply much more robust than many have realized.

One might worry that Thomason conditionals like (9b) and (9c) appear to conflict with the Ramsey test. If I hypothetically pretend to learn that my mother got me a sweater, then under that hypothesis, looking in the drawer is moot. The Ramsey test was meant to be a piece of evidence, beyond bare intuition, for the truth of the Miners Puzzle conditionals. So, the objection might go, Thomason

¹¹. *Objection*: If I look in the drawer and there's no necklace, I gain evidence that she didn't get me the necklace. *Reply*: Yes, but there's still some chance that she got me the necklace but didn't hide it in the drawer. So it's not the case that this action will probably lead to me knowing what she got me (it'll only lead to me knowing what she probably got me). So it doesn't conflict with my preference.

conditionals are in fact nonstandard, requiring a different analysis. The Miners Puzzle conditionals may still be analyzed as cases of ellipsis.

Reply: there are two interpretations of the Ramsey test.

If two people are arguing ‘If p will q ?’ and are both in doubt as to p , they are adding p hypothetically to their stock of knowledge and arguing on that basis about q ; so that in a sense ‘If p , q ’ and ‘If p , not q ’ are contradictories. We can say that they are fixing their degrees of belief in q , given p . (1931)

On the first interpretation, the conversational participants hypothetically pretend that p is part of their stock of knowledge. On the second, they merely add p to their actual stock of knowledge, generating a body of knowledge that isn’t necessarily theirs. This objection relies on the first interpretation.

The final quoted sentence is evidence that Ramsey intended the second interpretation: $Pr(q \mid p)$ need not be equivalent to $Pr(q \mid p \wedge Kp)$. The characterization of the role of information in conditionals that I defended in Section 1.3 also requires the second interpretation of the Ramsey test. The behavior of conditionals with information-sensitive operators like *probably* in their consequents generally seems to require the second interpretation. And the fact that there is a unified analysis of “standard” indicative conditionals, Thomason conditionals, and conditionals with information-sensitive operators is some evidence, though obviously not decisive evidence, that the analysis is correct. There is potentially more to say about nonstandard analyses of Thomason conditionals, but for the purposes of this paper I’ll assume that this is adequate evidence that the ellipsis account is not enough to explain the Miners Puzzle.^{12,13}

2. The Orthodox Account of Subjective *Ought* in Natural Language

2.1. Kratzer Semantics

Our semantics for indicatives shouldn’t validate certain classical inference rules, including at least modus tollens and proof by cases, for natural language indicative conditionals. This means trouble for some accounts of indicatives. The material conditional analysis of indicatives, of course, validates classical inference rules. So the behavior of information-sensitive operators in Miners Puzzle

12. Thanks to an anonymous referee for pressing me on this point.

13. There is another commonly suggested defense which I won’t discuss here: the suggestion that the *ought* takes wide scope over some kind of conditional and that for that reason modus ponens can’t be applied. Kolodny and MacFarlane (2010) and Silk (2014b), among others, provide strong evidence against this hypothesis, which I won’t rehearse.

cases provides strong evidence against the material conditional analysis.¹⁴

Furthermore, even accounts like Stalnaker (1975), which distinguish natural language indicative conditionals from material conditionals, validate classical inference rules. The reason is that the indicative conditional, on Stalnaker's view, entails the material conditional. Indicatives with information-sensitive expressions in their consequents show that this is false. Consider the example I gave in Section 1.3: (6b) does not entail (10).

(6b) If it's Monday, they're probably in shaft A.

(10) It's Monday $\supset$ they're probably in shaft A.

The material conditional analysis does not predict *probably* in (10) to be interpreted relative to a probability function that reflects conditionalization on *It's Monday*. But the unconditional probability of *they're in shaft A* is only 33%. So the consequent of (10) is false, but the antecedent may well be true. So the truth of (6b) does not entail the truth of (10).

Since the indicative conditional doesn't entail the material conditional, indicatives shouldn't be expected to validate classical inference rules just because the material conditional does. So the Miners Puzzle conditionals generate counterexamples to the Stalnaker semantics for indicatives, along with any other theory that makes indicatives entail material conditionals.

Fortunately, we're not left without an account of indicatives and subjective *ought*. In natural language semantics, the validity of modus ponens is not broadly accepted. The near-orthodox account of conditionals in ordinary language is Kratzer's (1981; 1991) restrictor analysis. And as Charlow (2013) and Yalcin (2012) observe, that analysis doesn't validate modus ponens.

Kratzer's restrictor analysis for conditionals treats them as a form of modalized sentence. On the restrictor analysis, all indicative conditionals contain a modal operator, which might be implicit. The antecedent restricts the domain of that operator.

Kratzer's semantics for modals relativizes them to two contextual parameters: a *modal base*, or set of propositions characterizing relevant background information about the world, and an *ordering source*, or set of propositions characterizing the contextually salient priorities or ideals. I'll call the intersection of propositions in the modal base the "modal background." We can think of the modal background as the set of epistemically or circumstantially possible worlds.¹⁵ The

14. The standard move made by defenders of the material conditional account is to claim that the truth conditions and assertability conditions of indicatives come apart, for pragmatic reasons. But at least the Gricean pragmatic account won't explain why the Miners Puzzle conditionals are assertable but not true, since that account only explains why some conditionals are true but not assertable (i.e., because they generate false implicatures).

15. While Kratzer's official position is that deontic modals take circumstantial modal

ordering source determines a partial ordering over worlds: better worlds are (roughly) worlds where more ordering source propositions are satisfied. These two parameters determine the *domain* of the modal: roughly, the best worlds within the modal background. Because it will be helpful later to use diagrammatic representations, I include the diagram format illustrated in Figure 1: the outer rectangle represents the modal background, broken up by the ordering; the domain of the modal is highlighted in gray.

Figure 1

Let a modal background i be a set of (for our purposes, epistemically) possible worlds. Let the ordering source d determine a deontic partial ordering in terms of ideality over those worlds.

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

In Kratzer's semantics, w is at least as good as w' iff the set of ordering source propositions that w satisfies includes the set of ordering source propositions that w' satisfies.

Definition 2. $O_{w,i,d}$ is the set of worlds in i such that $\leq_d$ doesn't rank any world in i higher.¹⁶ We can call $O_{w,i,d}$ the **domain** of the modal.

The linguistic orthodoxy is built from pairing the classic account of modals with Kratzer's account of conditionals. Where ' $\square$ ' is read *ought, should, must*, etc.,¹⁷ and $[[\cdot]]^{w,i,d}$ is the valuation function: a function from expressions in a language to extensions, and in particular, from sentences to sets of points (here, $\langle w, i, d \rangle$ triples where the sentence is true):

Modals: $[[\square\phi]]^{w,i,d} = \text{true}$ iff ϕ is true at all worlds in $O_{w,i,d}$.

Conditionals: $[[\text{if } \phi, \square\psi]]^{w,i,d} = \text{true}$ iff $[[\square\psi]]^{w,i \cap \phi, d} = \text{true}$.

bases, she doesn't restrict what it is for a modal background to be circumstantial. The only requirement is that it include propositions that characterize "relevant circumstances." One source of relevant circumstances may be agents' knowledge; and so epistemic modal bases are one form of circumstantial modal base.

16. I presuppose the limit assumption for convenience.

17. Kratzer's account doesn't distinguish strong and weak necessity modals.

2.2. Problems for Kratzer Semantics

I argued in Section 1 that no account of the Miners Puzzle that attempts to retain classical inference rules like modus tollens and proof by cases will prove satisfactory—at least not while saving the linguistic phenomena. But fortunately, the orthodox account of conditionals in semantics doesn’t validate these classical inference rules. So is the problem already solved?

No. As Kolodny and MacFarlane (2010) note, and Charlow (2013) and Cariani et al. (2013) argue at length, the orthodox semantics falsely predicts that relative to any given resolution of the contextual parameters the sentences in the original Miners Puzzle must be inconsistent. This is a manifestation of a deeper problem for that account: it inadvertently blocks certain kinds of interactions between information and the subjective *ought*.

First, let me rehearse Kolodny and MacFarlane’s argument for how Kratzer semantics goes wrong on the Miners Puzzle. The priority in this case is, I assume, to save as many of the miners’ lives as possible. So suppose our ordering simply orders worlds by how many of the ten miners are saved. Then the outer rectangle in Figure 2 represents the modal background *i*: here, the set of epistemically possible worlds. Our ordering source simply orders the worlds within the modal base according to how many of the miners’ lives are saved. And so on this ordering, the very best worlds—the domain $O_{w,i,d}$, which is again highlighted in gray—are the worlds where the speaker lucks out and blocks whichever shaft the miners happen to be in.

	In A	In B
Block A	10	0
Block B	0	10
Block neither	9	9

Figure 2: Priority: save as many miners as possible.

But if this is the ordering, then (1a) cannot be true:

(1a) We ought to block neither shaft.

So certainly this ordering doesn’t get us the right result. With this setting of the contextual parameters, Kratzer semantics predicts that the speaker should block one of the shafts.

It might be thought that what went wrong is that the ordering we used was an objective (information-independent) ordering. What we need is a subjective ordering, one that takes into account that we don’t know where the miners are.

So, perhaps a better ordering is one that orders worlds according to the *expected* number of miners’ lives that we save. And as Figure 3 shows, this allows us to get the right result for (1a): the worlds where we maximize the expected number of miners’ lives are all worlds where we block neither shaft.

	In A	In B
Block A	5	5
Block B	5	5
Block neither	9	9

Figure 3: Priority: maximize expected miners’ lives.

Unfortunately, Kratzer semantics can only get the right prediction for (1a) at the cost of getting the wrong prediction for (1b) and (1c).

- (1b) If the miners are in shaft A, we ought to block shaft A.
- (1c) If the miners are in shaft B, we ought to block shaft B.

We are now forced to predict that both are false, relative to this ordering. Figure 4 illustrates why. Consider (1b): the antecedent eliminates all of the worlds where the miners aren’t in shaft A. But because the ordering is fixed independently of the modal background, its ranking of worlds has to remain unchanged. And so since there are still worlds in the modal background where we block neither shaft, those are predicted to be the best worlds in the new, restricted modal background. But that’s clearly wrong.

	In A	In B
Block A	5	
Block B	5	
Block neither	9	

Figure 4: Priority: maximize expected miners’ lives.

But, it might be protested, conditional on the miners’ being in shaft A, surely the action with the greatest expected number of miners’ lives is blocking shaft A. And that’s obviously true. The problem is that in order to represent this, you need the ordering of worlds to be shiftable with changes in the modal background—which Kratzer semantics doesn’t allow. Kratzer semantics lets embedding under conditionals affect the modal background, but not the ordering.

Examples like the Miners Puzzle show that embedding under conditionals can shift the ordering as well.¹⁸

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

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

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

Cariani et al. (2013) offer a tidy general proof that Kratzer semantics can't allow serious information dependence: suppose for reductio serious information dependence. So there's some $w' \in i \cap \llbracket \phi \rrbracket$, such that (i) $w' \in O_{w,i,d}$ but (ii) $w' \notin O_{w,i \cap \llbracket \phi \rrbracket, d}$. Because of (ii) and Definition 1, there's some $w'' \in i \cap \llbracket \phi \rrbracket$ such that $w'' \leq_d w'$ but $w' \not\leq_d w''$. But then $w'' \in i \supseteq i \cap \llbracket \phi \rrbracket$. And so $w' \notin O_{w,i,d}$.

As a result, in order for Kratzer semantics to yield the correct prediction in this case, the priorities between (1a) and (1b) would have to differ. In other words, the sentences would have to exhibit some sort of equivocation. This is a bad result. There is simply no reason to predict that the priorities should change between the two assertions. Modal parameters aren't supposed to be infinitely flexible formal tools. Deontic ordering sources are supposed to be determined by contextually salient bodies of priorities. In the Miners Puzzle case, there is no evidence that these priorities change between utterances (1a) and (1b). If we allow for appeals to shifts in the parameter where there's no evidence of shifts in priorities, our metasemantics is going to be too flexible to be predictive or explanatory.

There is, by contrast, strong evidence that the subjective deontic *ought* shouldn't be expected to behave as the Kratzer account suggests. I think we can give a simple and decisive argument that subjective *ought* and other deontic modals are seriously information dependent.

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

Our semantics for deontic modals should allow for the possibility that, in some contexts, the salient priority is to maximize the expectation of some kind of value (money, hedons, miners' lives, etc.). So, for example, there's a reading in some context where *John should ϕ* is true iff *ϕ ing maximizes expected x* is true. On this reading, any world (in the modal base) where John maximizes expected x is a world where John is doing as he should.

The reading we're after makes *should*, like expected value, vary with a body of information. Importantly, that body of information can be picked out by the modal background parameter—the parameter that relativizes what's ideal to a set of relevant circumstances or information.

All of this should be uncontroversial. Now, a piece of data: this reading of the deontic necessity modal allows for the consistency of *John should ϕ* and *If χ , John should not ϕ , even though ϕ ing is still an open possibility*. It follows from the consistency of *ϕ ing maximizes expected value* and *if χ , then not ϕ ing maximizes expected value, even though ϕ ing is still an open possibility*. And so this is enough to generate serious information dependence.¹⁹

Note that the suggestion here is not that an adequate semantic account should incorporate a disambiguation of *ought* such that *ought ϕ* is true iff ϕ maximizes expected utility (as, for example, Lassiter, 2011, argues). The suggestion is rather that the semantics for deontic modals shouldn't rule out the possibility that maximizing expected utility is the sole salient priority in some contexts (e.g., in the casino with economist friends). We shouldn't build decision theory into our semantics, but we also shouldn't make the semantics incompatible with expressing the consequences of a decision theory. And that's what Kratzer semantics effectively does.

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

3. Some Conservative Fixes

Cariani et al. (2013) and Charlow (2013) independently give conservative amendments to the orthodox semantics that involve similar basic operations: they introduce a third parameter that, under certain circumstances, coarsens the ordering over worlds. In the Kratzer semantics, the ordering source ordered worlds. These accounts order options. Doing so allows them to predict the consistency

19. This point is also made by Cariani (in press).

and truth of the Miners Puzzle sentences.

Cariani et al.'s third parameter is a "decision problem": a partition over the modal background such that each cell of the partition represents a different option that is choosable to the agent. An action α is *choosable* iff there's some action specification β such that it's epistemically necessary that the agent can knowingly perform β and it's epistemically necessary that performing β entails that α is achieved. So, for example, in the Miners Puzzle scenario, the action of saving nine miners is choosable. Saving ten miners is not choosable: there's no action such that we can knowingly perform that action and that doing so will save ten miners.

As with Kratzer, the ordering is determined by a set of propositions. But whereas Kratzer defined the ordering in a fine-grained way, over worlds, Cariani et al. coarsen the ordering so that worlds are only ranked differently if they are in different cells of the partition of options.²⁰

In other words, even if an option can have many possible outcomes with different objective values, the ordering of worlds doesn't reflect those differences: it only orders *options*, not outcomes. Why? We don't *know* what outcomes each option will bring about. An option α is ranked better than an option β if the set of ordering source propositions α entails includes the set of ordering source propositions β entails.

In the Miners Puzzle case, for example, suppose the ordering source is {we save ten miners, we save nine miners, ..., we save one miner}. Now, what we should do is be in a cell that entails as many of these propositions as possible. But since the cells are determined by which actions are choosable, there's no cell that represents the action of saving all ten miners. The partition, instead, will be roughly {we block shaft A, we block shaft B, we block neither shaft} (as in Figure 5). The cells where we block one of the shafts include worlds where we save all miners and worlds where we kill all miners. Since they include the latter, they entail none of the ordering source propositions. But the cell where we block neither shaft entails that we save nine miners, and so entails nine of the ordering source propositions.

20. Formally, Kratzer's orderings are determined as follows, where d is an ordering source (a set of propositions determining the partial ordering $\leq_d$):

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

whereas Cariani et al.'s orderings are determined as follows, where $[w]_\pi$ is the cell (option) of a partition π that a world w is in:

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

	In A or B
Block A	0
Block B	0
Block neither	9

Figure 5: Guaranteed miners’ lives saved.

So the best option according to our ordering is to block neither shaft, and so that option is the domain of the modal (again, highlighted in gray). So the Cariani et al. account makes the correct prediction for (1a).

When the modal background is restricted to worlds where the miners are in shaft A (as in Figure 6), the cell where we block shaft A includes only worlds where we save all ten miners, and so entails ten ordering source propositions. Blocking neither shaft only entails nine. So Cariani et al. also make the correct prediction for (1b).

	In A	In B
Block A	10	
Block B	0	
Block neither	9	

Figure 6: Guaranteed miners’ lives saved with a restriction to A-worlds.

Charlow’s (2013) account offers somewhat different motivations but generates the same results. Charlow views his account as an application of the account of weak necessity modals in von Fintel and Iatridou (2008). On their account, weak deontic necessity modals quantify over all of the best worlds according to a secondary ordering within the set of best worlds according to the primary ordering. Strong deontic necessity modals quantify over all best worlds according to a primary ordering.

Charlow argues that the explanation for the Miners Puzzle sentences derives from the interaction of a secondary ordering source with a primary ordering source like the one characterized above. Whereas the primary ordering source requires that we save as many miners as possible, the secondary ordering source requires that we perform actions that bring about ends that are *actionable* and avoid those that are not, where actionability is basically the same as choosability. Now, the secondary ordering is used to coarsen the primary orderings in the same way as the decision problem in Cariani et al., and so in the Miners Puz-

zle scenario the two accounts will make the same predictions about sentences involving weak deontic necessity modals.²¹

4. The Deeper Problem with the Orthodoxy

Kratzer, in unpublished notes, argues that Cariani et al.'s account, however conservative, involves some unnecessary addition of decision theoretic machinery, in particular, the decision problem parameter. Kratzer asks rhetorically, "Why pack information about rational decision making into the meaning of modals?"—the implicature being, of course, that we shouldn't.

I agree with Kratzer that we shouldn't pack information about decision theory into the meaning of modals. But the Kratzer account, the Cariani et al. account, and the Charlow account all do. And they don't pack in just any information: rather, they include substantive structural normative assumptions, in the form of controversial decision theoretic rules. In Section 4.1, I'll develop some problem cases for all three accounts and then in Section 4.2, I'll explain how these cases are manifestations of normative commitments that are built into the orthodox assumptions about modals and conditionals. For that reason, conservative amendments to the Kratzerian orthodoxy will not yield an adequate account of subjective *ought* and other deontic modals. We need a more radical reconceptualization of how information and priorities interact in normative language.

4.1. Problems with Accounting for Probability and Value Differences

Consider a variation on the Miners Puzzle scenario:

21. The cost of tying his account to the von Fintel and Iatridou account is that Charlow's account only produces reasonable predictions for weak deontic modals; he incorrectly predicts that (i) is false:

- (i) We mustn't block either shaft.

Von Fintel (2012) also rejects this judgment. Charlow takes it to be an advantage for his view that it uses tools from von Fintel and Iatridou and so has independent theoretical motivation. But Charlow's secondary ordering sources formally behave quite differently from von Fintel and Iatridou's. Their secondary ordering sources make finer grained distinctions among the worlds that are best according to the primary ordering. His secondary orderings remove distinctions within the primary ordering, yielding a coarser grained ordering. It's not much of an exaggeration to say that the only similarity seems to be that both accounts replace Kratzer's single ordering source with two moving parts. For example, Charlow also predicts both of the following sentences are true in the Miners Puzzle context (and are therefore consistent):

- (ii) a. We should block neither shaft.
b. We must block one of the shafts.

So on Charlow's account, weak deontic modals are not weaker than strong deontic modals.

High probability. Everything is the same as in the original Miners Puzzle, except that instead of having no idea where the miners are, we're 99.9% confident that they're in shaft A. Nothing else is changed: priorities are held fixed, and the same options are available; only our information has changed. There's a very good chance we can save all of the miners.

(11) is true:

(11) We ought to block shaft A.

Cariani et al. and Charlow are unable to predict (11) (relative to the fixed set of priorities). The problem: their account coarsens the ordering over worlds so that it only distinguishes between chooseable or actionable outcomes: that is, outcomes that we can (knowingly) bring about. So they predict *ought* to be insensitive to the probabilities of outcomes given our actions.

On Cariani et al.'s account, any world w where we block shaft A is in the same cell of the decision problem partition as a world w' where we block shaft A and the miners are in shaft B. After all, that's still an epistemic possibility for us. w' is a world where we save no miners. So the option of blocking shaft A does not entail any of the propositions in the ordering source. But the option of blocking neither shaft still entails saving nine lives (i.e., entails nine ordering source propositions). And so by their ranking, any world where we block neither shaft is better ranked than any world where we block shaft A; so *We ought to block neither shaft* is still predicted to be true.

Charlow's account has the same effect. On his view, there is a normative demand on us (at least in some contexts) to perform actions that we know will bring about certain of our ends. In the probabilistic Miners Puzzle case—where, again, we hold fixed all features of the original Miners Puzzle case except the salient probabilistic information—this means that we must perform the action that we know will save some of the miners' lives, and that we must not perform the action that we are reasonably confident *but not certain* will save all of the miners' lives.

Kratzer's account exhibits the same insensitivity to adjustments of probability. In our example, the modal background still contains worlds where the miners are in shaft A and worlds where they're in shaft B. (The latter are just less probable.) And so worlds where they're in A and we block A are no better than worlds where they're in B and we block B. So (11) is predicted to be false: some ideal worlds are block-B worlds. It doesn't matter how improbable they are.

Other sorts of counterexamples emerge: consider a case where there's a 95% chance the miners are in shaft A and otherwise they are in neither mineshaft. There's no downside whatsoever to blocking shaft A. Still, these accounts predict that *We should block shaft A* is false. In Cariani et al.'s account, for example,

the worlds where we block shaft A and save all the miners' lives are in the same cell as the worlds where we block shaft A but don't save anyone by doing so. There are always worlds in that cell where we do no better (vindicate no more propositions in the ordering source) than if we did nothing. So that cell is neither better nor worse than the cell where we do nothing. The result: while *We shouldn't block shaft A* is false, so is *We should block shaft A*.

The coarsening accounts, like Kratzer semantics, are also not sensitive to cardinal differences in the value (desirability, moral status, etc.) of outcomes. These accounts all only consider the ordering, not whether one outcome is *a lot* better than another. Consider a variant of the original Miners Puzzle:

One Miner. Everything is the same as the original Miners Puzzle case, except that blocking neither shaft will only save one of the miners, instead of nine.

The coarsening account predicts that (1a) will be true.

(1a) We should block neither shaft.

While the option of blocking a shaft doesn't entail any of the ordering source propositions, blocking neither shaft entails one of them. So Cariani et al. predict (1a) to be true. Blocking neither shaft brings about one of our actionable ends, while blocking a shaft doesn't, so Charlow also predicts (1a). Finally, if Kratzer were to predict the truth of (1a) in the original Miners Puzzle case, relative to these priorities, she would have to predict it here too, at least on any natural ordering source.

But (1a) is very plausibly false. In that circumstance, the best thing to do might be to block a shaft at random.²²

Of course, all three accounts could easily devise ad hoc ordering sources that happened to give the correct predictions for individual cases. But for any alternative ordering source that they suggest, we can devise counterexamples that have the same structures. Cardinal differences in probability and value matter for some of our normative judgments. For this reason, they matter for some of our *ought* claims. This doesn't mean that we have to build cardinal probabilities into our semantics. But it does mean, at minimum, that a semantic theory should be able to reproduce the predictions of a cardinally enriched account without artificially stipulating changes in deontic contextual parameters.²³

22. If this judgment is unclear, change the scenario so that there are 1,000 miners' lives at stake and blocking neither shaft will still only save one.

23. Compare: Holliday and Icard (2013) argue that an adequate semantics for comparative epistemic modals need not include probability functions or other cardinal information about likelihood; comparative probability orderings may be all that's needed. But Holliday and

4.2. Decision Theoretic Commitments Built into the Semantics

These empirical objections to the Kratzer, Cariani et al., and Charlow accounts point to a more general objection. Both accounts incorporate serious normative assumptions into their semantics. Moreover, because these normative assumptions aren't the sorts of assumptions that guide most ordinary language speakers' judgments, they lead to a variety of false predictions.

Cariani et al. rule out the truth of $\Box\phi$ if and only if the ϕ -worlds in the modal background include even one world with an objectively worse outcome than the worst outcome possible for some alternative to ϕ . Normatively, this amounts to the assumption that an action α is worse than an action β if and only if the worst possible outcome of α is worse than the worst possible outcome of β .

In other words, Cariani et al. encode the decision rule Maximin into the semantics of deontic modals.²⁴ This is not a good thing to build into the semantics of deontic modals.

So I agree with Kratzer that their account builds decision theoretic information into the meanings of normative language.²⁵ But in fact, Kratzer's account is guilty of the same charge. While the Cariani et al. account encodes the decision rule Maximin, Kratzer's account encodes Maximax: roughly, the rule that one should choose the option that has some positive probability of yielding the best possible outcome. This is a straightforward consequence of the more naive commitment of that semantics: that we should always simply bring about the best possible outcome in the modal background.²⁶

Icard's account is able to capture all the predictions of the stronger account. By contrast, Cariani et al. can only make the predictions I suggest by stipulating changes in the deontic contextual parameter resolution.

24. Charlow's semantics, which I earlier mentioned faced many of the same objections as Cariani et al. account, avoids this charge on a technicality. (His normative view entails Maximin, but it isn't semantically encoded.) But his account will still yield the same problematic predictions as the Cariani et al. account, in the same range of cases. Moreover, his semantics does still encode some form of normative commitment, in that it rules out certain decision rules—in particular, decision rules that are sensitive to probabilistic information and cardinal value differences.

25. However, Kratzer suggests that the partition of options in Cariani et al.'s account is already problematically decision theoretic, and I don't. I take issue instead with their semantics' decision theoretic commitments.

26. Cariani (in press), partly in response to a manuscript of this paper, claims that one cannot tell from structural features of an account's lexical entries for deontic modals, whether the account encodes a decision rule. Cariani provides no argument for this claim, and it seems to me false. Decision rules like Maximin are themselves structural: they are merely functions from orderings over outcomes and sets of alternatives to selected alternatives. The orderings Maximin uses need not be *utility* orderings, construed as ordering by subjective desirability. The relevant value might be miners' lives, or the number of miners' lives greater than seven, or quantity of beer in the fridge, etc. The orderings could even be functions of probabilistic information. For example, Γ -Maximin selects alternatives on the basis of maximizing minimum

This observation leads to my two main objections. First, the semantics of modals shouldn't build in substantive decision theoretic commitments. Second, they certainly shouldn't encode implausible decision theoretic commitments.

The commitments to Maximax or Maximin in the two accounts I discussed can't be written off as a byproduct of some idealization. Neither of the views under discussion allows the subjective *ought* to be sensitive to probabilities. And while decision rules like Maximax and Maximin show no sensitivity to probabilities, other decision rules do.

More generally, we want our semantics for subjective deontic modals to at least allow for sensitivity to probabilistic information and cardinal differences in value, because as a matter of normative fact, our subjective obligations are sensitive to both of these factors. And as a matter of psychological fact, our judgments about what we ought to do are affected by these factors. Plausibly, any view that makes the information-sensitive operator *probably* insensitive to probabilities of some sort couldn't really be considered an account, no matter how idealized, of the natural language *probably*.²⁷ It seems to me that the same is true of the information-sensitive *ought*: if a theory doesn't allow (not force, just allow) *ought* to show sensitivity to probabilities or cardinal differences in value, it's not a theory of the subjective *ought*.

However, we should also reject theories, like Lassiter's (2011) account of Miners Puzzle phenomena, which yields desirable predictions only by explicitly encoding Bayesian decision norms. While we might find such norms more decision-theoretically attractive than Maximin or Maximax, it's unreasonable to predict that ordinary agents' deontic, deliberative *ought* assertions always express claims

expected value, for cases where there are multiple candidate probability assignments. Of course, conflicting decision rules can yield the same verdicts if they're fed different orderings, particularly if these orders are not made to be held fixed while other features of an example are tweaked. But as a structural matter, if we hold fixed an ordering over outcomes and a set of alternatives in a context, then relative to that ordering, Cariani et al. (2013) necessarily predicts that the true deontic *should/may* utterances in that context will coincide with the verdicts of Maximin for that ordering (if Maximin yields any verdicts at all).

One may reply that the problem for Cariani et al. (2013) and others isn't major, because the verdicts for Maximin, relative to one ordering, could coincide with the verdicts of other decision rules, relative to different orderings. But in order for the (Cariani et al. 2013) semantics to achieve the verdicts that seem plausible at a context, where, e.g., the speaker's intended decision rule is expected miners' lives maximization, we need to gerrymander the ordering. Orderings that achieve this result will not coincide with what would naturally be understood as the speaker's goals in the Miners Puzzle, i.e., saving miners' lives. For example, one can reach the right verdict for (1a) in the *One Miner* case if, e.g., the ordering doesn't distinguish saving one miner's life or none. But in the natural construal of the case, with sane priorities, that ordering doesn't coincide with the speaker's values. The upshot: we can force the right predictions only by making the ordering both less natural and less transparent to the speaker and to theorists. This too seems to me a substantial cost.

27. But see Holliday and Icard (2013).

that accord perfectly with Bayesianism. For example, agents might have tacit or explicit commitments to other sorts of (arguably rational) decision rules, like risk-weighted expected utility theory (Buchak 2013), sunk-cost sensitive decision theories (Doody 2013), E-admissibility, Γ -Maximin (Seidenfeld 2004), etc. Should we bar agents from expressing these commitments in *ought* claims?

It's important to avoid building normative assumptions into natural language semantics. As Lewis (1978) wrote: "The semantic analysis tells us what is true (at a world) under an ordering. It modestly declines to choose the proper ordering. That is work for a moralist, not a semanticist." The main reason: if you hold a belief that is inconsistent with some normative claim and you express your belief linguistically, your claim might be false. But it shouldn't show a failure of linguistic competence! If controversial decision theoretic norms are built into semantics, that means one can't even coherently *debate* the deliverances of these norms. And as Moore's Open Question Argument suggests, there may be no norms whatsoever that have that status: no norms that are part of the very meaning of simple normative expressions like *ought*.

5. The Proposal

5.1. Predictiveness without Normative Commitments?

It is a cost to a theory if it makes semantics take sides in normative theorizing. Kolodny and MacFarlane's (2010) view doesn't have this problem. On their view, a contextually determined deontic selection function takes a subset of an information state (modal background) as the domain for quantification of deontic modals. This view is a generalization of Kratzer's account and its offshoots.

But unlike those other accounts, the account in Kolodny and MacFarlane (2010) is not a predictive account. Kolodny and MacFarlane give no explanation of how the deontic selection function works, what semantic or contextual factors determine which deontic selection function is in play in some context, etc. So their account doesn't seem to give enough materials to make any sort of predictions of why the Miners Puzzle sentences, or any other deontically modalized sentences, are true; it only allows us to predict that they are consistent. (See Charlow, 2013, and Cariani et al., 2013, for a fuller development of this objection.)

This lack of predictive power suggests that the view isn't fully explanatory. It's hard to base an explanation on a primitive selection function. For comparison: suppose Rae is the tallest woman in Maryland. We might ask: why is Rae so tall? It doesn't seem like a full answer to say: "Well, there's a selection function from geographic regions to people that picks out the tallest person in each region, and Rae is the value of that selection function applied to Massachusetts."

In order to provide a more predictive account, we need to be able to say something more. An analogy: suppose we have two accounts of the indexical *I*. The first account says, simply, that *I* is context sensitive; its referent is a function of the context. The second account fills in how context is relevant, in a way that allows the account to be predictive: it tells us, for example, that unembedded uses of *I* refer to the *speaker* in the context, perhaps also that the referent of *I* can shift when it appears in certain kinds of embeddings (Santorio 2012), etc. It's easy to see how the second story has an advantage in predicting which sentences containing *I* will come out true. A more predictive account of how deontic modals behave will similarly tell us what kinds of facts about a context will determine the interpretation of the modal, give a sense of how different kinds of embeddings can affect the interpretation of modals, and so on.

There are therefore (at least) three desiderata that extant accounts of subjective *ought* and other deontic modals don't jointly satisfy:

1. The account should allow for serious information dependence.
2. The account should not build in substantive normative commitments.
3. The account should be predictive.

One might reasonably ask whether this is a rigged game: how can we predict the truth of deontically modalized sentences without having substantive normative commitments in the semantics? Trivially, if we predict the truth of *it ought to be the case that ϕ* , then we are committed to the fact that it ought to be the case that ϕ . So it seems like it should be impossible to satisfy desiderata 2 and 3.

But it is not impossible to satisfy both desiderata. Normative commitments can arise in the linguistic account without arising within the semantics. Compare pronouns: nothing in the semantics of the word *he* allows us to predict that on some occasion, when I use *he*, I'll be referring to Richard.

Consider the classic account of modals: with modals that are not information-sensitive—for example, objective *ought*—the classic account manages to be reasonably predictive without incorporating normative assumptions. We can think of the classic account of deontic modals as incorporating two covert pronouns for the two contextual parameters.²⁸ One pronoun picks out something like a body of information (which determines the modal background) while the other picks out some set of priorities (which determines an ordering of possibilities). Perhaps, like other pronouns, they can act as free or bound variables. *According to Massachusetts state law, we should...* could be an example of a bound occurrence. When they're not bound, the resolution of the parameters is determined by the contextually salient priorities or bodies of information, just as the reference of *he* is determined by contextually salient males.

28. I offer this suggestion for expository purposes; my account is not committed to this hypothesis about the contextual parameters.

The ability of the classic account to generate predictions hinges on the possibility of avoiding ad hoc stipulations about the ordering source. For this reason, I think it is a serious theoretical cost that Kratzer semantics cannot predict that the Miners Puzzle sentences are all literally true and share the same ordering source. Intuitively, no changes in the salient priorities takes place between the Miners Puzzle sentences:

- (1a) We ought to block neither shaft.
- (1b) If the miners are in shaft A, we ought to block shaft A.

Since no unique ordering can be plugged into the classic Kratzer semantics to predict their mutual truth in the Miners Puzzle scenario, Kratzer semantics should be revised. Stipulated changes in the ordering between the two sentences, or predictions that one or the other of the sentences is literally false, handle the Miners Puzzle only at the cost of making the pragmatic side of the Kratzer account too flexible to be predictive. Still though, we should take our cue from Kratzer in letting context supply norms, instead of hard-coding them into the semantics.

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

5.2. *Two Normative Parameters*

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

What kinds of decision rules does r look for? There are plenty of alternatives: most obviously, expected value maximization, but there could also be Maximin, Maximax, risk-weighted expected utility maximization, rational sunk cost reasoning, or whatever. In all cases, the kind of value in question is provided by the priorities parameter, analogous to the ordering source.

This means we have two normative parameters. Why not just some sort of function from information states to sets of worlds where the subjectively best

actions are performed? This is the route that Kolodny and MacFarlane (2010) take. But they do nothing to explain what, in a particular case, would determine the relevant function. I think the most plausible way of spelling out what determines a deontic selection function in a context will be a story of the form I'm characterizing: a story according to which, in a particular context, the domain of quantification for deontic modals is determined by a function from what's valued, what's uncertain, and some decision rule to what's subjectively best.

There are also benefits in using two separate normative parameters. While one allows for serious information dependence, the other should not. And having a parameter that does not allow for serious information dependence and reflects objective sources of value means that we can easily recover the objective *ought* within the same semantic account. How? One option is by giving the decision-making parameter an uninteresting resolution, one that says simply to bring about the best possible outcome in the modal base. So, for any possible source of value, we can represent the close relation between the subjective and objective *ought* for that sort of value: they share the same value function and only differ with respect to the decision rule parameter. Another option is to allow for resolutions of *r* that ignore the priorities parameter *v* and so mimic traditional ordering sources.

Kolodny and MacFarlane (2010) argue that there is no objective *ought* in natural language. After all, they claim, this sort of usage is useless for the purposes of deliberation. But this is false: even in deliberative contexts, the objective *ought* has important uses. For example, in the Miners Puzzle case, the following sentences are only true on the objective reading of *ought*:

- (12) a. For all I know, we ought to block shaft A.
 b. It might be that we should block shaft A.

Moreover, there are clearly some interpretations of *ought* that demand an objective interpretation: the legal *ought*, for example, is not sensitive to anyone's knowledge or ignorance. So we should acknowledge the existence of objective as well as subjective *oughts*. The kind of account I want to offer makes clear how they are related.

5.3. *The Machinery*

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

First, I argued above that subjective deontic modals need to show sensitivity to probabilities of some sort. A modal that doesn't allow for sensitivity to probabilities is simply not the subjective *ought*. So, just as we need to access some sort of probability function for the information-sensitive operator *probably*, the

same will be true for subjective *ought*. Second, I argued that we should make our priorities-sensitive parameter potentially sensitive not just to ordinal ranking of different outcomes, but also to cardinal differences between outcomes. These two changes mean that we can no longer represent either of those two parameters as mere sets of propositions.

So, our three parameters:

1. An information state s . Following Yalcin's (2010) suggestion for the semantics of *probably*, I'll represent an information state as a pair $\langle i, Pr \rangle$ of a modal background (a set of worlds) and a probability function.²⁹
2. A value function v . The obvious way to represent this is as a function from worlds to real numbers.^{30,31}
3. A decision rule parameter r . r is a set of decision norms that determine an ordering over worlds in the modal background.

For simplicity, I treat s and v as precise. This isn't a commitment of the account; we could instead, for example, treat s as sensitive to sets of probability functions, or to credence functions that don't obey the probability axioms, or sets of such credence functions, or whatever; similarly for v .³²

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

30. This is an oversimplification. If we just use a unique value function, we lose one advantage that Kratzerian ordering sources have: the ability to represent incommensurability. I stick with a unique value function for the sake of simplicity of exposition, but a fully fleshed out version of this account will need *sets* of value functions in order to fulfill two roles: (i) representing incommensurability and (ii) allowing for sensitivity to cardinal differences in value.

31. *Objection*: Doesn't your view build in some kind of normative assumptions? Isn't this effectively building in consequentialism? *Reply*: While it is a vexed question whether all normative theories can be consequentialized, nothing I say here suggests that they can. First, the ranking of worlds doesn't necessarily depend on *consequences*. A world might be better if someone is rational, or has a good heart, or is brave, or whatever. Second, the semantics says nothing about whether value should be maximized. Third, assigning worlds values doesn't require rejecting agent-centered norms: v doesn't need to be a ranking of worlds in terms of moral value in general. It can just be a ranking of worlds in terms of agent-centered value.

32. Rothschild (2012) argues that in the case of the *probably* operator, we actually need sets of probability functions in order to handle a puzzle about disjunctions of *probably* sentences. The puzzle is that there are true disjunctions of the form *probably* $\phi \vee$ *probably* ψ where neither disjunct is true. I am skeptical that this adequately motivates moving from probability functions

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

Modals: $[\Box\phi]^{w,s,v,r} = \text{true}$ iff ϕ is true at all worlds the modal's domain, $O_{w,s,v,r}$

Conditionals: $[\text{if } \phi, \Box\psi]^{w,s,v,r} = \text{true}$ iff $[\Box\psi]^{w,s+\phi,v,r} = \text{true}$

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

The semantics for modals is simply a generalization of Kratzer's semantics that allows modals' domain to be affected by the new parameter r . Similarly for the semantics for conditionals, which also updates the salient probability function(s) by conditionalization, in the same way that Yalcin (2010) argues is necessary for probability operators.³³

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

5.4. The Decision Rule Parameter

The rational decision parameter, r , is a generalization of a Kratzerian ordering source. Ordering sources, recall, were sets of propositions. They were able to determine an ordering of worlds: w is d -better than w' iff w satisfies a strict superset of the d propositions that w' satisfies.

Instead of sets of propositions, we can use sets of functions from $\langle s, v \rangle$ pairs to propositions. Call a set of such functions an r -ordering source. This amounts to an interpretation of the kinds of things that decision-theoretic norms are: they are norms that don't merely assess worlds, but instead assess worlds relative to bodies of information and priorities. We saw that ordering sources containing sets of propositions like those in (13) couldn't yield the appropriate kind of information-sensitivity:

- (13) a. {We save 10 miners, we save 9 miners, ..., we save 1 miner}
b. {We maximize expected value}

to sets of probability functions. It seems to me that these are cases where the disjunction's *probably*s reflect some form of objective probability where one or the other disjunct is true (but the speaker doesn't know which). Neither disjunct is true only relative to the speaker's subjective probabilities or credences.

33. Note that this new semantics abstracts away from complications, noted by Yalcin (2007) and Kolodny and MacFarlane (2010), about the behavior of conditionals with information-sensitive modals in their antecedents. I'm happy to take on board their suggested solution.

The amendments to the classic semantics that I have argued for—in effect, making the ordering source potentially sensitive to cardinally structured information and priorities—are neutral about the question of whether modals are assessment-sensitive.

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

- (14) a. {We perform an action that maximizes s -expected v -value}
 b. {We perform an action that maximizes the minimum s -possible v -value}
 c. {We perform an action that maximizes the maximum s -possible v -value}³⁴

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

- (15) {We don't violate any federal laws $\langle s, v \rangle$ }

where (15) is a constant function of $\langle s, v \rangle$.

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

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

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

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

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

Contrast this account with Kratzer's in the treatment of the Miners Puzzle case. Suppose we use the Kratzerian ordering source {We maximize the expected number of miners' lives}. Worlds where we maximize expected miners' lives—relative to our subjective probabilities—are worlds where we block neither shaft. But what about (1b)?

- (1b) If the miners are in shaft A, we ought to block shaft A.

The antecedent of the condition—which is merely a piece of linguistic material—doesn't affect our information. We don't learn anything from the *if*-clause. Our information is just determined by the evidence we've received, the facts that we know. The body of information that is shifted by the antecedent isn't our information. It is, instead, a somewhat more abstract entity, a body of linguistically

34. There is no bar to r -ordering sources containing multiple decision rules. Note that if some decision theoretic norms are, in a context, prioritized over other decision theoretic norms, we might need sequences of r -ordering sources, as in von Fintel and Iatridou (2008).

useful information: the body of information determined by a combination of our information and any restrictions that could be added on by any *if*-clause imaginable. This body of information, *s*, can be changed by surrounding linguistic material. Hence, the clause *if the miners are in shaft A* gives us an information state that is more informed than ours. The *s*-expectation, unlike our expectation, can update on the information in the *if*-clause. And so an *r*-ordering source like {We maximize *s*-expected *v*}, where *v* is the number of miners' lives saved, can correctly predict (1b) to be true in exactly the same context where *We ought to block neither shaft* is true.

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

5.5. Upshot

The account I've given is another generalization of Kratzer's semantics for modals and conditionals. The Cariani et al. and Charlow accounts are still too strong: they constrain the choice of *r* to specific kinds of decision rules. The Kolodny and MacFarlane account, meanwhile, is too weak: it fails to be predictive. My account aims to find a happy middle.

I noted above that there are (at least) three desiderata for an account of the subjective *ought* and other information-sensitive normative expressions.

- (i) The account should allow for serious information dependence.
- (ii) The account should not build in substantive normative commitments.
- (iii) The account should be predictive.

Kratzer's account satisfies only (iii): her account rules out serious information dependence and builds in Maximax. The Cariani et al. account satisfies (i) and (iii), but it still builds in Maximin. The Kolodny and MacFarlane account satisfies (i) and (ii), but doesn't tell us how to predict the truth or falsity of *ought* claims.

My account is designed to satisfy all three desiderata. First, it allows for serious information dependence: the *r*-generated ordering of worlds is partly a function of the information state. Second, my account avoids normative commit-

ments: these are left to be resolved by context. Third, my account is at least as predictive as Kratzer's or Cariani et al.'s. It tells you something about the inner workings of the semantics, and where to go looking for the resolutions of the contextual parameters. And with those, we can start making predictions. For example: if v ranks worlds according to how many of the miners live, and r is maximizing the s -expectation of v -value, then we make the right predictions for the Miners Puzzle sentences.

6. Conclusion

I've described some surprising behavior exhibited by the subjective *ought* and other deontic modals. First, subjective *ought* claims can violate classic inference rules. Second, subjective *ought* claims exhibit serious information dependence. I've argued that one can give an account of subjective *ought* and other information-sensitive modals that predicts and explains their behavior without any substantive normative commitments. Other accounts that have been offered either fail to be predictive or incorporate unwarranted normative assumptions into the meanings of modals. The account I provide respects the motivations for the classic account of modals and conditionals, but excises the normative commitments latent in the orthodox account.

References

- Bledin, Justin (2014). Logic Informed. *Mind*, 123(490), 277–316.
- Buchak, Lara (2013). *Risk and Rationality*. Oxford University Press.
- Cariani, Fabrizio (in press). Deontic Modals and Probabilities: One Theory to Rule Them All? In Nathan Charlow and Matthew Chrisman (Eds.), *Deontic Modality*. Oxford University Press.
- Cariani, Fabrizio, Magdalena Kaufmann, and Stefan Kaufmann (2013). Deliberative Modality under Epistemic Uncertainty. *Linguistics and Philosophy*, 36(3), 225–259.
- Charlow, Nate (2013). What We Know and What to Do. *Synthese*, 190(12), 2291–2323.
- Doody, Ryan (2013). The Sunk Cost “Fallacy” is Not a Fallacy. Manuscript. Retrieved from <http://www.mit.edu/~rdoody/TheSunkCostFallacy.pdf>
- Dowell, Janice L. (2012). Contextualist Solutions to Three Puzzles About Practical Conditionals. In Russ Shafer-Landau (Ed.), *Oxford Studies in Metaethics* (Vol. 7, 271–303). Oxford.
- Holliday, Wesley H. and Thomas F. Icard (2013). Measure Semantics and Qualitative Semantics for Epistemic Modals. *Proceedings of SALT 23*, (514–534).

- Kolodny, Niko and John MacFarlane (2010). Ifs and Oughts. *The Journal of Philosophy*, 107(3), 115–143.
- Kratzer, Angelika (1981). The Notional Category of Modality. In Hans-Jürgen Eikmeyer and Hannes Rieser (Eds.), *Words, Worlds, and Contexts: New Approaches in World Semantics* (38–74). de Gruyter.
- Kratzer, Angelika (1991). Modality. In Arnim von Stechow and Dieter Wunderlich (Eds.), *Semantics: An International Handbook of Contemporary Research* (639–650). de Gruyter.
- Lassiter, Daniel (2011). *Measurement and Modality: the Scalar Basis of Modal Semantics*. (Doctoral dissertation). New York University.
- Lewis, David (1978). Reply to McMichael. *Analysis*, 38(2), 85.
- Parfit, Derek (1988). What We Together Do. Manuscript. Retrieved from http://individual.utoronto.ca/stafforini/parfit/parfit-_what_we_together_do.pdf
- Ramsey, Frank (1931). General Propositions and Causality. In Richard B. Braithwaite (Ed.), *The Foundations of Mathematics and Other Logical Essays* (235–55). Kegan Paul, Trench, Trubner & Co.
- Regan, Donald (1980). *Utilitarianism and Cooperation*. Oxford.
- Rothschild, Daniel (2012). Expressing Credences. *Proceedings of the Aristotelian Society*, 112(1pt1), 99–114.
- Santorio, Paolo (2012). Reference and Monstrosity. *Philosophical Review*, 121(3), 359–406.
- Seidenfeld, Theodore (2004). A Contrast Between Two Decision Rules for Use with (Convex) Sets of Probabilities: Γ -Maximin Versus E-Admissibility. *Synthese*, 140(1-2), 69–88.
- Silk, Alex (2014a). Evidence Sensitivity in Weak Necessity Deontic Modals. *Journal of Philosophical Logic*, 43(4), 691–723.
- Silk, Alex (2014b). Why ‘Ought’ Detaches: Or, Why You Ought to Get with My Friends (If You Want to Be My Lover). *Philosophers’ Imprint*, 14(7), 1–16.
- Stalnaker, Robert (1975). Indicative Conditionals. *Philosophia*, 5(3), 269–286.
- van Fraassen, Bas (1980). Review of Brian Ellis, *Rational Belief Systems*. *Canadian Journal of Philosophy*, 10(3), 497–511.
- von Fintel, Kai (2012). The Best We Can (Expect to) Get? Challenges to the Classic Semantics for Deontic Modals. Manuscript. Retrieved from <http://web.mit.edu/fintel/fintel-2012-apa-ought.pdf>
- von Fintel, Kai and Sabine Iatridou (2008). How to Say Ought in Foreign: the Composition of Weak Necessity Modals. In Jacqueline Guéron and Jacqueline Lecarme (Eds.), *Time and Modality: Studies in Natural Language and Linguistic Theory* (Vol. 75, 115–141). Springer.
- Yalcin, Seth (2007). Epistemic Modals. *Mind*, 116(464), 983–1026.
- Yalcin, Seth (2010). Probability Operators. *Philosophy Compass*, 5(11), 916–37.

Yalcin, Seth (2012). A Counterexample to Modus Tollens. *Journal of Philosophy of Logic*, 41, 1001–1024.