

Deliberative Modality under Epistemic Uncertainty

Fabrizio Cariani, Magda Kaufmann & Stefan Kaufmann

May 14, 2013

1 Introduction

The modalized sentence in (1) is easily understood as expressing that reading said article is the thing to do.

(1) We ought to read that article.

It is standard practice in linguistic semantics to treat a modal like *ought* as expressing quantification over possible worlds that are selected and ranked according to certain contextual parameters (Kratzer, 1977, 1981, 1991, 2012). On this view, different modal flavors correspond to different values for these contextual parameters. In the following, we will call the particular modal flavor at play in practical deliberations like (1) *deliberative modality*,¹ and we will assume that, together with at least deontic (in the narrow sense, that is, related to rules and laws) and bouletic (related to wishes) it belongs to the so-called *prioritizing* modal flavors (in contrast to epistemic and dynamic ones, cf. Portner 2009).²

The ability of this approach to provide an adequate account of inferences involving sentences that express deliberative modality has been questioned. One recent such challenge is raised by the discussion of the *miners puzzle* in Kolodny and MacFarlane (2010):³

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.

¹See Thomason (1981) for a philosophical characterization of the concept of deliberative modality; for a recent, non-equivalent, list of criteria also Schroeder (2010).

²Deliberative modality may be closely related to teleological, that is, goal oriented, modality. However, teleological modality is sometimes claimed to involve an additional ‘in order to’-parameter (a *designated goal*, cf. von Fintel and Iatridou ms.). Since we do not want to go into this discussion, and no such parameter is needed for the cases of practical deliberation we are concerned with, we keep the distinction between deliberative modality (which does not require a designated goal) and teleological modality (which may require a designated goal).

³The puzzle was originally discussed, in a different theoretical context, by Parfit (1988, 2011), who credits it to Donald Reagan. For an analogous case, see also Jackson (1991).

Action	if miners in A	if miners in B
Block shaft A	All saved	All drowned
Block shaft B	All drowned	All saved
Block neither shaft	One drowned	One drowned

(Kolodny and MacFarlane 2010, p. 155-ff)

Kolodny and MacFarlane argue that this puzzle poses a logical and a semantic problem. As we will show, the semantic problem directly affects the standard linguistic picture of deliberative modality. For the situation described, they report the following judgments:⁴

- | | | | |
|-----|----|--|------|
| (2) | a. | We ought to block neither shaft. | TRUE |
| | b. | If the miners are in shaft A, we ought to block shaft A. | TRUE |
| | c. | If the miners are in shaft B, we ought to block shaft B. | TRUE |

These judgments strike us as correct (at least on one reading). In principle, it would be best to block the shaft in which the miners are, thus (2-b) and (2-c) seem to be true. But given that we don't know and have no way of finding out, it is best to block neither shaft and thus guarantee that we save nine of the ten. So, (2-a) seems to be true as well.⁵ However, from (3), which is undisputedly true in the scenario under consideration, together with (2-b) and (2-c) under standard assumptions of propositional logic (disjunction introduction, disjunction elimination, and modus ponens for indicative conditionals) we can derive (4).

- | | | |
|-----|--|-------|
| (3) | Either the miners are in shaft A or they are in shaft B. | TRUE |
| (4) | Either we ought to block shaft A or we ought to block shaft B. | FALSE |

But (4) is intuitively incompatible with (2-a). And indeed, (4) doesn't seem to be a correct outcome of practical reasoning.

Before we discuss in more detail the theoretical implications raised by these data, let us briefly clarify two issues arising from additional readings of the sentences under consideration.

First, the deliberative interpretation of *ought* that underlies the judgments discussed most prominently in Kolodny and MacFarlane's paper is not the only one available, as they acknowledge and we agree. (2-a) also has an *objective reading* on which it is not felt to be true. Objectively, it would be best to block the shaft the miners are actually in; consequently, on this objective reading (4) feels true. While we will occasionally get back to the objective reading when comparing the predictions

⁴Like Kolodny and MacFarlane, we consider only readings of (2-a) on which the negative existential takes narrow scope with respect to *ought*. As far as we can tell, the inverse (non-surface) scope order does not display any phenomena relevant to our discussion over and above what is discussed in connection with the other examples we consider.

⁵It is vital to Kolodny and MacFarlane's (2010) reasoning that the decision has to be taken without there being a way of finding out where the miners are. Otherwise we would get different judgments on (2-a): given that human lives are at stake, speakers broadly agree that the agents ought to do their best to find out where the miners are, rather than make do with the limited information they have. Knowing where the miners are, they could save all ten for sure, while with the limited knowledge they have they can only save nine for sure, or even risk the death of all ten.

resulting from particular theoretical choices, we generally focus on the problem of how to predict the deliberative reading.

Second, the above conditionals also have an additional reading on which they are not intuitively true. This reading can be brought out as follows:

- (5) 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.

We call this reading *non-reflecting*. Once again, our emphasis is on the reading on which (2-b) and (2-c) are true in the scenario (the *reflecting* reading), but the additional reading in (5) will have a role to play in our discussion (§2.3.3).

Focusing for now on the deliberative reading, the miners puzzle brings out two distinct problems:

LOGICAL: How should we block the derivation of (4) from (2-b), (2-c) and (3)?

SEMANTIC: How can our semantic machinery derive all of the salient verdicts?

Kolodny and MacFarlane (2010) discuss strategies to solve the logical problem by rejecting one or more of the premises or by giving up on one or more of the inference rules. They end up jettisoning modus ponens; without this rule, they argue, from (2-b), (2-c) and (3) we can no longer derive (4). As they acknowledge, this leaves open the question concerning the correct semantics.⁶

In much of the recent literature on natural language semantics, neither the validity of modus ponens for indicative conditionals nor a classical analysis of disjunction is taken for granted.⁷ In light of this, we might expect that this literature would hold a ready-made solution to the semantic problem. This is not the case, however. In §2, we show that the problem arises given any straightforward specification of the parameters in the standard framework.⁸ In §3, we propose an extension of the standard framework and show how it leads to an account of the semantic problem.

Our solution to the miners puzzle shares a central feature of Kolodny and MacFarlane's solution to the semantic problem (*serious information dependence*, cf. §3). At the same time, we argue that it improves on their proposed answer to the semantic problem in three respects: First, it gives a better diagnosis of the puzzle. Second, it offers a more systematic account of how the interpretation of deliberative modals

⁶A word of caution is in order at this point: Clearly the logical problem arises only if the reading of *ought* stays constant across (2-a), (2-b) and (2-c). Kolodny and MacFarlane argue explicitly that it does. But the truth of the undisputably deliberative (5) calls into question whether *ought* in (2-b) and (2-c) is indeed deliberative as well, rather than objective. This issue is hard to settle in the absence of independent evidence, but since our focus is on the semantic side of the problem, not much hinges on the logical side for our purposes.

⁷For some examples of the first kind, see Lycan (1993, 2001); Cantwell (2008). For recent discussions of disjunction, see Zimmermann (2000); Geurts (2005); Simons (2007); Barker (2010).

⁸Charlow (forthcoming) argues for the same negative conclusion. Although related, our arguments are somewhat different and complementary. Charlow does not consider certain escape routes which we leave open (in particular what we call the Covert Operator Construal in §2.3.3). Charlow also offers a different resolution to the problem than we do. We defer discussion of his approach to future work. Charlow also discusses different English necessity modals and their propensity for objective and deliberative readings, respectively.

changes as context evolves. In §4, we discuss the relation of our proposed extension to the standard framework. We show that it is a conservative extension: For all non-deliberative cases, our approach can simulate the standard framework (§4.1). Furthermore, we show that for deliberative modals in unembedded contexts, our proposal can be mimicked in the standard framework (§4.2). Nevertheless, we argue (§4.3) that this implementation does not extend to embedded occurrences, and sacrifices some of the intuitive appeal of Kratzer’s analysis. In that sense, we consider our proposal, which technically expands on the standard framework, as more faithful to the spirit of Kratzer’s framework and more adequate to the data (since it handles both the embedded and the unembedded uses).

2 The miners problem in Kratzer’s framework

In this section, we briefly introduce the standard Kratzer-style framework and argue that it does not offer a solution to the miners problem. We conclude the section by arguing that it is nonetheless a useful guide in looking for a solution.

2.1 Formal background

Let a non-empty set W of possible worlds be given. Propositions are represented as subsets of W . Atomic sentences denote propositions. We first define the interpretation of modal auxiliaries like *ought* and then discuss various ways of applying the framework to the sentences in the miners example.

Following Kratzer (1981, 2012), modal sentences are interpreted relative to two parameters, usually labeled f and g . Both are *conversational backgrounds* in her terms – formally, they are functions from possible worlds to sets of propositions. The first is the *modal base*. It determines which worlds form the domain relative to which the modal expressions in question are to be interpreted, much like the accessibility relations familiar in modal logic. For each world w , we call the set $\bigcap f(w)$ (the worlds at which all propositions in $f(w)$ are true) the *modal background* at w .⁹ The second parameter g is the *ordering source*. For each world w , the propositions in $g(w)$ are used in the comparative evaluation of possible worlds according to their stereotypicality, desirability, compliance with norms, or other criteria, depending on the modal flavor embodied in g . Technically, Kratzer defines a pre-order $\leq_{g(w)}$ on the set of worlds, where ‘ $u \leq v$ ’ may mean that u is less far-fetched than v , preferable to v , or “better” than v in some other respect:

$$(6) \quad u \leq_{g(w)} v \text{ iff } \{p \in g(w) \mid v \in p\} \subseteq \{p \in g(w) \mid u \in p\}$$

When $u \leq_{g(w)} v$ but not $v \leq_{g(w)} u$, we write ‘ $u <_{g(w)} v$ ’. In the interpretation of modals, the ordering source is used to select the most “relevant” worlds (for the

⁹There is some variation in the literature as to whether the term “modal base” for a given world w refers to the function f itself, to $f(w)$ (the value of this conversational background), or to $\bigcap f(w)$ (the set of accessible worlds in the sense of modal logic). We use the term as defined in Kratzer (1981), namely for the conversational background. If there is no danger of confusion, we also use it for the value $f(w)$ at a particular world w . However, we reserve the term “modal background” for the set of worlds $\bigcap f(w)$.

given modal flavor) from the modal background: these are the *minimal* worlds under the induced ordering. For instance, in the case of a deontic or bouletic ordering source g , the propositions in $g(w)$ indicate what is required or desirable. The more of these propositions are true at a given world u (i.e., the closer u is to the “ideal”), the lower it is ranked. The minimal worlds (i.e., the ones closest to the “ideal”) are the ones relevant for the interpretation of deontic or bouletic necessity and possibility statements.

If $g(w)$ is finite, there is guaranteed to be a set of minimal worlds in $\bigcap f(w)$. Since nothing in the given scenario turns on the question whether the ordering source is finite, we assume without loss of generality that it is.¹⁰

We refer to the set of minimal worlds as in (7).

$$(7) \quad O(f, g, w) := \left\{ u \in \bigcap f(w) \mid \neg \exists v \in \bigcap f(w) [v <_{g(w)} u] \right\}$$

We further stipulate that in the cases of interest in this paper the modal base is always **consistent** – formally, $\bigcap f(w)$ is non-empty for all f, w . Together with the preceding constraint, this implies that $O(f, g, w)$ exists and is non-empty.

Our assignments of truth-conditions are relativized to triples of the form $\langle f, g, w \rangle$ (which we may call *points of evaluation*). Later, as we add a third conversational background, we will want to streamline this notation: for this reason, in §3.1, we will bundle the conversational backgrounds into a single parameter ‘ c ’ ranging range over *contexts*.

Based on this formal apparatus, we now define a notion of **necessity** that is parametrized to both the modal base and the ordering source and defined in terms of quantification over the minimal worlds in the modal background. Formally, this is represented using the modal operator \boxplus in the object language, whose semantics is as spelled out in (8).^{11,12}

$$(8) \quad \boxplus\varphi \text{ is true at } \langle f, g, w \rangle \text{ iff for all } v \in O(f, g, w), \varphi \text{ is true at } \langle f, g, v \rangle.$$

A special case of the human necessity operator is the (familiar) **simple necessity** operator \square , defined in (9) as shorthand for human necessity relative to a constant ordering source whose value is the empty set (of propositions). In this case, the truth of the sentence depends only on the modal base and the world of evaluation.

$$(9) \quad \square\varphi \text{ is true at } \langle f, g, w \rangle \text{ iff } \boxplus\varphi \text{ is true at } \langle f, \lambda v.\emptyset, w \rangle.$$

¹⁰ The assumption that every linearly ordered chain within the partial order terminates in a set of minimal worlds is called the **Limit Assumption** after Lewis (1973). It does not imply that $g(w)$ is finite, but the converse holds. Nothing in our account hinges on either the finiteness of $g(w)$ or the Limit Assumption, but we adopt both for ease of exposition.

¹¹This type of quantificational force is called “human necessity” in Kratzer (1981). The ‘H’ in the operator \boxplus is mnemonic for that. Kratzer (1991, 2012) calls it simply “necessity.” On either usage it is to be distinguished from the special case of “simple necessity” defined in (9).

¹²Without the Limit Assumption (see Fn. 10), $O(f, g, w)$ may be undefined. The definition of necessity in (8) could then be rewritten along the following lines (cf. Kratzer, 1981, 2012):

(i) $\boxplus\varphi$ is true at $\langle f, g, w \rangle$ iff for all $v \in \bigcap f(w)$, there is some $v' \in \bigcap f(w)$ such that (i) $v' \leq_{g(w)} v$ and (ii) for all $v'' \in \bigcap f(w)$ such that $v'' \leq_{g(w)} v'$, φ is true at $\langle f, g, v'' \rangle$.

The overarching goal of the Kratzerian enterprise in modal semantics is to capture the semantic variability and context-dependence of modal expressions in terms of the modal base and ordering source. According to the received view, *ought* expresses necessity relative to particular settings of these parameters. What are those settings in the miners case?

2.2 Ordering source

The standard assumption about the ordering source for *ought* is that it represents goals or moral obligations. In principle, these two notions are distinct and give rise to *teleological* and *deontic* modal flavors, respectively. In the given scenario, intuitions may waver as to how the particular ordering source should be labeled. But assuming that one’s moral obligations are also goals one strives to attain, this question can be set aside for our present purposes. What is important is that we can identify a stock of propositions that govern our decision as to what step to take next, and they obviously include “that as few miners as possible die.”¹³

We refer to this ordering source as *g*. Crucially, its value does not depend on what shaft the miners are in, or indeed on any other fact left unspecified by the description of the scenario. Thus *g* is constant across all relevant worlds compatible with what we are told about the scenario. In order to account for the intuition that a world in which we save $n + 1$ miners comes out better than one in which we save only n , we assume that at each world w compatible with what we know, the ordering source returns the set of propositions given in (10).¹⁴

$$(10) \quad g(w) = \{ \text{all miners are saved, at least 9 miners are saved, } \dots, \\ \text{at least 1 miner is saved} \}$$

2.3 Modal base

The choice of the modal base is more complicated than that of the ordering source. In the following, we discuss the main issues from the perspective of a Kratzer-style

¹³Note that this may follow from a more general principle like “if someone is in danger, save that person.” As there are no other lives at stake that are relevant under the current circumstances, we can ignore more general moral obligations that do not immediately pertain to the deliberation described in the scenario. For example, the deliberation seems unaffected by the possibility to give up our jobs, study medicine and join *Médecins Sans Frontières*, which might enable us to save even more (totally unrelated) lives. Zooming in on the particular case in that sense seems to correctly capture speakers’ intuitions about the sentences in (2-a)-(2-c).

¹⁴Angelika Kratzer (p.c.) and Kai von Stechow (ms.) questioned our assumption that these are the relevant preferences in the scenario as described and proposed alternative accounts of the data. Both note that the prospects for an orthodox analysis improve if we use modalized propositions in the ordering source. A full-fledged comparison between our proposal and Kratzer’s will have to wait until her approach is worked out in complete detail, but at the time of this writing, we believe that there are reasons to prefer our account. In the case of von Stechow’s proposal, we are unsatisfied with the prediction that the conditionals are false when *ought* gets its deliberative interpretation. Von Stechow (pp. 28-30) maintains that the truth judgment can be recovered by showing that some *different but related* conditionals involving deliberative *ought* are true. We believe (although there naturally is room for further argument) that the datapoint is that it is *these* (i.e. (2-b)-(2-c)) very conditionals that are true on the deliberative interpretation of *ought*.

Table 1: Classes of worlds in the given scenario

<i>cell</i>	<i>position of miners</i>	<i>action</i>	<i>miners lost</i>
AA	shaft A	block A	0
AB	shaft A	block B	10
AN	shaft A	block neither	1
BA	shaft B	block A	10
BB	shaft B	block B	0
BN	shaft B	block neither	1

semantics, which we will later compare to Kolodny and MacFarlane’s proposal.

Before we start, notice that the set of possible worlds compatible with the description of the scenario can be split into six mutually exclusive subsets, each of which is characterized by a state of affairs (i.e., location of the miners) and an action on the part of the narrators (‘we’). For ease of discussion, we refer to these six possibilities with the letter sequences given in the leftmost column in Table 1. For each possibility, the two parameters jointly determine the outcome (i.e., the number of miners lost) listed in the rightmost column. With this picture in mind, we turn to two kinds of modal bases that may be involved in the interpretation of *ought* and consider their respective predictions about the miners scenario.

2.3.1 Circumstantial modal base

In the linguistic literature, the modal base for *ought* is typically taken to be *circumstantial* – that is, determined by the relevant facts of the situation. The question of what counts as a “relevant” fact is difficult to answer in general, but intuitions are clear enough for a simple scenario like the miners problem.¹⁵ What is important about the notion of a circumstantial modal base is that it is independent of – though not necessarily distinct from – the information available to the deliberating agents involved in the situation. In this respect, it comes closest to what Kolodny and MacFarlane discuss as the *objective* interpretation of the modal. We refer to this modal base as f_o .

The description of the scenario does not give the reader all the relevant facts that would be required to determine the content of the circumstantial modal base. Clearly the location of the miners is of vital importance for the consequences of our actions, and the description of the situation makes it clear that the matter is settled objectively: There is a shaft, A or B, such that the miners are in it. By any plausible criterion for relevance in the given scenario, this is a relevant fact, and so the circumstantial modal base includes it. But we (both the ‘we’ of Kolodny and MacFarlane’s description and we as readers of their paper) do not know which of the two shafts it is.

Formally, the possible worlds that are consistent with the description of the scenario fall into two classes according to whether the miners are trapped in shaft A or B. In Table 1, they are indicated as the sets $AA \cup AB \cup AN$ and $BA \cup BB \cup$

¹⁵For some general reflections on this issue, see §3.4.

Table 2: Parameters of evaluation with circumstantial modal base

	w_A	w_B
modal background	$\bigcap f_o(w_A) = AA \cup AB \cup AN$	$\bigcap f_o(w_B) = BA \cup BB \cup BN$
ordering source	$g(w_A) = g(w) = g(w_B)$	
ranking ¹⁶	$AA \cup BB \leq_{g(w)} AN \cup BN$	$AN \cup BN \leq_{g(w)} AB \cup BA$
minimal worlds	$O(f_o, g, w_A) = AA$	$O(f_o, g, w_B) = BB$

BN, respectively. Let w_A and w_B be two arbitrary worlds from each of these two classes. The world w in which the described agents carry out their deliberation is similar to one of them in all relevant respects. Thus the value of the circumstantial modal base f_o at w is one of the two sets of propositions in (11). The description does not specify which one it is.

- (11) $f_o(w) \in \{f_o(w_A), f_o(w_B)\}$
- a. $f_o(w_A) = \{\text{there are two shafts A and B,}$
10 miners are in shaft A,
 we can't find out if the miners are in A or in B,
 blocking a shaft rescues everyone in that shaft,
 either one shaft is blocked or both shafts will fill to a level
 so that a person at its bottom dies}
- b. $f_o(w_B) = \{\text{there are two shafts A and B,}$
10 miners are in shaft B,
 we can't find out if the miners are in A or in B,
 blocking a shaft rescues everyone in that shaft,
 either one shaft is blocked or both shafts will fill to a level
 so that a person at its bottom dies}

The parameter settings relevant for the evaluation of (2-a) at either of these worlds are as given in Table 2. Recall that the sentence is true under these settings at a world w if and only if all worlds in $O(f_o, g, w)$ are such that we block neither shaft. But this is neither the case at w_A nor at w_B : In each of the corresponding modal bases, the minimal worlds are the ones at which we block the shaft that the miners are in. Thus, even though we don't know exactly what the world is like, we do know that (2-a) is false on the circumstantial reading. In contrast, (4) is true at both w_A and w_B .¹⁷

- (2-a) We ought to block neither shaft. ✗
- (4) Either we ought to block A or we ought to block B. ✓

¹⁶Expressions of the form ' $X\rho Y$ ', where X, Y are subsets of the domain and range of the binary relation ρ , are shorthand for the statement that $x\rho y$ for all $x \in X, y \in Y$.

¹⁷Here and in the following, ✗ and ✓ indicate the predictions of the respective theory under discussion, rather than speaker intuitions.

These predictions reflect the objective reading for *ought* (see §1), but they are at odds with the deliberative reading we are interested in, i.e., the reading that plays a crucial role in the agents' practical reasoning.

As an aside, we actually think that the above judgments reflect an interpretation of the sentences that is not only available to us as the readers of the paper, but even to the deliberating subjects themselves. Another example with similar status is given in (12) (where the existential quantifier takes scope over the modal).

(12) There is a shaft we ought to block. ✓

Here, too, our intuition is that while (12) may be judged false in view of the deliberative process (in particular, it is at odds with (2-a)), it can be understood as true on the objective interpretation. The combination of a circumstantial modal base and a deontic ordering source predicts (12) to be true.

Putting that issue aside, we grant that the reading Kolodny and MacFarlane have in mind (on which (2-a) is true and (4) is false) is also real and not captured by the combination of a circumstantial modal base with a deontic ordering source. In order to account for this deliberative reading, we need to keep looking for the right combination of parameters.¹⁸

The problem with the circumstantial modal base is that it includes the actual location of the miners.¹⁹ This information is not available to the deliberating agents, thus useless in their deliberations. Perhaps, then, a better approach is to take into account only the circumstances that are in fact epistemically accessible to the agents.

¹⁸ One point about the deliberative/objective distinction deserves mention here. Note that in order for a sequence like (i) to be consistent, the first two occurrences of *ought* must receive an objective interpretation, while the third occurrence has to be deliberative.

- (i) We ought to block shaft A or we ought to block shaft B, but the problem is that we don't know which one. In view of that limitation, we ought to block neither.

We do think that (i) is consistent, and the shift required to explain this fits well with the standard view in linguistics that modal verbs are context dependent and obtain different readings depending on what conversational backgrounds are salient. Kolodny and MacFarlane (2010) explicitly reject solutions to the miners problem that postulate different readings for the occurrences of *ought* in (2-a)-(4). It is important to note, however, that this objection is orthogonal to our concerns. Although we believe that (2-a) and (4) can be used – with the appropriate context shift – in one and the same argument, our solution to the miners problem does not depend on such a shift. Rather, our point is that the truth of (2-a) in the given situation cannot be captured by the standardly assumed settings for modal base and ordering source.

¹⁹ Angelika Kratzer (p.c.) objected to this claim. However, we are not aware of a general definition of the notion of a circumstantial modal base that would justify the exclusion of the miners' location in the given scenario. Nor do we consider our claim to be at odds with Kratzer's (2012) proposal to abandon the dichotomy of epistemic vs. circumstantial modal bases in favor of the more encompassing notion of *realistic* ones (i.e., conversational backgrounds f such that $w \in \bigcap f(w)$ for any $w \in W$): Any realistic modal base that excludes the miners' location yields the same problematic predictions as the epistemic modal base considered in the following section. Thus the issue turns out to be irrelevant to our argument.

Table 3: Classes of worlds in the epistemic modal base

	w
modal background	$\bigcap f_e(w) = AA \cup AB \cup AN \cup BA \cup BB \cup BN$
ordering source	$g(w)$
ranking	$AA \cup BB \leq_{g(w)} AN \cup BN \leq_{g(w)} AB \cup BA$
minimal words	$O(f_e, g, w) = AA \cup BB$

2.3.2 Epistemic modal base

The foregoing suggests an *epistemic* modal base representing the information at the agents' disposal. This information is in principle independent of the facts, but we do assume for simplicity that it is *truthful*, i.e., not in conflict with the relevant facts – in our framework, with the content of the circumstantial modal base.²⁰ The crucial feature of the available information is that it can (but need not) be *incomplete*, as it is in Kolodny and MacFarlane's scenario. For the purposes of our discussion of the miners scenario, we call a modal base *uninformed* if it does not specify the location of the miners, and *informed* if it does.²¹

The uninformed case. The information available to the agents in Kolodny and MacFarlane's scenario is represented by a modal base f_e whose value at the world w of evaluation is specified by description as the set of propositions in (13). Notice that the value of f_e , unlike that of f_o above, is the same at all worlds consistent with the description.

- (13) $f_e(w) = f_e(w_A) = f_e(w_B)$
 = {there are two shafts A and B,
 10 miners are either in shaft A or in shaft B,
 blocking a shaft rescues everyone in that shaft,
 either one shaft is blocked or both shafts will fill to a level so that
 a person at its bottom dies}

The resulting parameter settings for the evaluation of the sentences in question are shown in Table 3. The set of minimal worlds under the induced ordering is $AA \cup BB$; at all these worlds, we block a shaft. So the sentence (2-a) is again not true. In this regard, the move from f_o to f_e did not help matters.

- (2-a) We ought to block neither shaft. ✗
 (4) Either we ought to block A or we ought to block B. ✗

²⁰The truthfulness assumption seems to be shared by Kolodny and MacFarlane. Kratzer (2012), in assuming that all potentially relevant modal bases are realistic (cf. Fn. 19), is committed to it. For us it is merely a matter of convenience, as it simplifies the comparison of subjective and objective readings of the *ought*-sentences.

²¹Notice that this distinction does not depend on the truthfulness of the information: A modal base which contains false information about the miners' location would nonetheless be informed in our sense.

However, the epistemic modal base does yield a different prediction about (4): Since now the set of minimal worlds comprises both ones in which we block A and ones in which we block B, neither of the disjuncts is true under this reading, and thus the disjunction is false. Similarly, since the minimal worlds are divided according to which of the two shafts is blocked, (12) is now predicted to be false: There is no shaft such that we block it at all minimal worlds.

(12) There is a shaft we ought to block. ✗

In sum, these predictions do not quite fit our intuitions about deliberative *ought*, under which (2-a) should be true. Nor do they capture the objectivist judgment that (2-a) is false whereas (4) and (12) are true. Recall that these were captured correctly by the circumstantial modal base. Thus the epistemic modal base fails to deliver a satisfactory Kratzer-style analysis for either of the readings of *ought*-sentences that can be observed in the miners scenario.

The informed case. At this point one might think that we just happened to pick the wrong epistemic modal base by relying on the one that specified the deliberating agents' information. In principle, an epistemic modal base might correspond to any other agent's (and possibly an even more abstract) body of information.²²

For concreteness, consider an alternative *informed* epistemic modal base f'_e whose values at the worlds compatible with the scenario are like those of f_e , except for the addition of truthful information about the location of the miners: Thus at each A-world w_A , \llbracket the miners are in A $\rrbracket \in f'_e(w_A)$; similarly for the B-worlds.²³ Now, without any further assumptions about f'_e , it is easy to see that in terms of the resulting predictions about our *ought*-sentences, f'_e behaves just like f_c . Specifically, at all A-worlds w_A , the modal background $\bigcap f'_e(w_A)$ contains only A-worlds, and the best ones among them are ones at which shaft A is blocked. Similarly, *mutatis mutandis*, for the B-worlds.

Stepping back, it is evident that any body of information consistent with the scenario leads to problematic predictions: If it specifies the location of the miners, it behaves like the circumstantial modal base in predicting (2-a) to be false because at all best worlds the shaft in which they are gets blocked. If it does not specify the location of the miners, it predicts (2-a) to be false because the best worlds are divided

²²This flexibility has recently been widely discussed in the literature on epistemic modals (DeRose, 1991; Egan et al., 2005; Egan, 2007; von Fintel and Gillies, 2007, 2008; Hacking, 1967; MacFarlane, 2011; Stephenson, 2007, among others).

Kolodny and MacFarlane (2010) argue that variability according to whose information state is relevant should not be captured as dependence on the context of utterance, but on the context of assessment. Our solution could be implemented under either of these assumptions, but we avoid further discussion of this issue here and stick with the more familiar notion of context dependence.

²³A comment is in order about the relationship between f_e and f'_e . In principle, we can think of the latter as derived from the former via an *update* with the information about the miners (see §2.3.3 below). On the other hand, in our view the central judgments about (2-a) and similar sentences rely on the agents' inability to find out where the miners are: If they had a way to find out, they ought to do so and then block the right shaft, rather than blocking neither. In order to reconcile this feature with the possibility of their finding out where the miners are, we would have to assume that the information comes to them unexpectedly, through a channel which they did not know or think to check themselves.

into worlds where A is blocked and worlds where B is blocked.

2.3.3 Deliberative *ought* in conditionals

The previous sections showed that Kratzer’s approach does not yield a straightforward account of the deliberative reading of the standalone *ought*-sentences in the miners scenario. We now turn to the conditionals (2-b) and (2-c).

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

In Kratzer’s formalization, conditional antecedents restrict the modal bases of modal operators. This can be seen as a reflection of Ramsey’s (1929) idea of hypothetically adding the antecedent to one’s stock of knowledge; however, as we are about to see, the Kratzerian framework offers some flexibility in spelling out the details. Before we discuss the relevant options, we introduce the original proposal of Kratzer (1981; 1986; 2012; and elsewhere).

We adopt Kolodny and MacFarlane’s notation and refer to the result of modifying a modal operator O with a conditional antecedent *if* φ as ‘ $[\text{IF } \varphi] O$ ’. For all modal bases f and propositions p , we refer to the result of an *update of f with p* as $[f+p]$: Formally, $[f+p]$ is a conversational background (i.e., a function from worlds to sets of propositions) like f , and such that for all worlds w , $[f+p](w) = f(w) \cup \{p\}$. It is easy to see that this results in a restriction of the modal background $\bigcap f(w)$ to worlds at which p is true.²⁴ Relative to this restricted modal base, the modified operator receives its usual interpretation. Illustrating with the operator \boxplus :

- (14) $[\text{IF } \varphi] \boxplus \psi$ is true at $\langle f, g, w \rangle$ iff $\boxplus \psi$ is true at $\langle [f+[\varphi]], g, w \rangle$
iff for all $v \in O([f+[\varphi]], g, w)$, ψ is true at $\langle [f+[\varphi]], g, v \rangle$.

It is generally assumed that when no overt modal is present in the sentence, a covert epistemic necessity modal is filled into the logical form and the antecedent serves to restrict its modal base. Once the option of a covert modal operator is admitted into the theory, two construals become available for all conditionals containing overt modals: Under the first construal, the overt modal is the main operator of the sentence and targeted by the modification with the *if*-clause. Under the second construal, the *if*-clause modifies the covert epistemic operator, which in turn has the overt modal in its scope (i.e., in effect embedded in the consequent of the conditional). M. Kaufmann, writing as Schwager (2006), labels these two interpretations the **Overt** and **Covert Conditional Operator** construal, respectively. The two are schematically given for overt deontic *ought* in (15) (‘ \boxplus_e ’ stands for the silent epistemic operator, ‘ \boxplus_d ’ represents the overt deontic modal).²⁵

²⁴Notice that $O(f, g, w)$ and $O([f+[\varphi]], g, w)$ stand in no particular relation. For instance, they are disjoint whenever both are non-empty and φ is false at all worlds in $O(f, g, w)$. This captures the non-monotonic behavior of conditionals, i.e., the fact that it is possible that ‘ $\boxplus \psi$ ’ is true while ‘ $[\text{IF } \varphi] \boxplus \psi$ ’ is false.

²⁵Here we use the subscripts ‘*e*’ and ‘*d*’ as an informal device to refer to epistemic and deontic interpretations, respectively, not to signify concrete parts of the logical form.

- (15) If φ , ought ψ .
- | | | |
|----|---|-----------------------------------|
| a. | $[\text{IF } \varphi] \boxplus_d \psi$ | Overt Conditional Operator (OCO) |
| b. | $[\text{IF } \varphi] \boxplus_e [\boxplus_d \psi]$ | Covert Conditional Operator (CCO) |

Instantiated with sentence (2-b), the two construals correspond to (16-a) and (16-b), respectively.

- (16) If the miners are in shaft A, we ought to close shaft A.
- | | | |
|----|---|-------|
| a. | $[\text{IF } [\text{the miners are in A}]] \boxplus_d [\text{we close A}]$ | (OCO) |
| b. | $[\text{IF } [\text{the miners are in A}]] \boxplus_e [\boxplus_d [\text{we close A}]]$ | (CCO) |

Kolodny and MacFarlane consider only the OCO construal for the conditionals they discuss. But in the linguistic literature, both construals have their advocates, and the current consensus is that the CCO construal is needed for at least some examples, especially ones involving deontic modals. Indeed, some authors argue that the CCO construal may well be correct for *all* conditionals. Although a detailed discussion of the relevant arguments would lead us too far afield,²⁶ the topic is obviously relevant because conditionals with (overt) deontic modals constitute a central data point in both discussions. We avoid taking a stance in this paper on which construal is correct so we need to consider both.

Overt Conditional Operator Construal. There is room for variation in the choice of modal base f – it may be circumstantial or epistemic, and in the latter case it may represent any number of information states or information sources – but regardless of what f represents, under the OCO construal the main clause $\boxplus_d [\text{we close A}]$ is interpreted at w relative to $\cap [f+A](w)$. (For simplicity, we use this as shorthand for ‘ $\cap [f+[\text{the miners are in A}]](w)$ ’.) Clearly this is a subset of the original modal background $\cap f(w)$, containing only worlds in which the miners are in shaft A. Assuming that the ordering source ranks worlds according to the number of miners saved, the best ones are those at which the agents block shaft A. Similarly for (2-c) and shaft B.²⁷ Thus the OCO construal easily accounts for the truth of the conditionals (2-b) and (2-c).

Covert Conditional Operator Construal. Under both construals, the consequent is evaluated at worlds in the modal background at which the antecedent is true. The two construals differ in that under CCO this evaluation of the consequent involves an additional modal operator (the one that is overtly present in the conditional sentence), complete with its own modal base and ordering source.²⁸ For ease of exposition, we

²⁶See Frank (1996), Zvolenszky (2002), Geurts (ms.), and Kaufmann and Schwager (2011) among others.

²⁷There are some “pathological” cases among the possible OCO-readings. For instance we argued above that a circumstantial modal base $f_o(w)$ contains the actual location of the miners at w . But then, the addition of the antecedent of (16-a) is either trivial or contradictory, as the case may be. We have nothing to say here about the usefulness of this reading or its pragmatic felicity.

²⁸There is no *a priori* reason to assume that the parameters of interpretation for the two modal operators stand in any particular relation. However, not all possibilities seem equally plausible in all examples.

dub the parameters of the covert epistemic operator f, g , and the ones for the overt deontic operator f', g' .

We assume with Kratzer that f is epistemic. Plausibly, it represents the uninformed epistemic state of the deliberating agents under uncertainty as to the location of the miners (see §2.3.2 above). We write f_e for this modal base. It could in principle be accompanied by an ordering source g of the kind typical for epistemic modals, i.e., one encoding likelihood or stereotypicality. However, since such notions do not play any role in the miners scenario, we assume that g has no effect, thus $O(f_e, g, w) = \bigcap f_e(w)$.

The consequent $\boxplus_d[\text{we block A}]$ is evaluated point-wise at all worlds in $\bigcap f_e(w)$ at which the antecedent is true, i.e. the miners are in shaft A. We assume as before that the ordering source g' for *ought* is deontic and ranks worlds according to the number of miners saved (we write g'_d); in addition, we assume that the agents are not uncertain about their preferences: for all worlds u, v such that $u, v \in \bigcap f_e(w)$, we have $g'_d(u) = g'_d(v)$. The choice of the modal base f' for *ought* requires some discussion.

CCO+o: Circumstantial (f'_o). The consequents in (2-b) and (5) are interpreted at the A-worlds in $\bigcap f_e(w)$. Recall from §2.3.1 above that under the circumstantial construal, $\boxplus_d[\text{we block A}]$ is true and $\boxplus_d[\text{we block neither}]$ is false at all A-worlds. The B-cases are similar.

- | | | |
|-------|--|--------|
| | | f'_o |
| (2-b) | If the miners are in A, we ought to block A. | ✓ |
| (5) | If the miners are in A, we ought to block neither. | ✗ |

CCO+e: Epistemic (f'_e). We are particularly interested in the case that f'_e is obtained from f_e , the modal base of the covert outer epistemic operator. There are two possibilities, each corresponding to an intuitively available reading of the conditional antecedent.

First, in (17) the modal base results from updating f_e with the antecedent.²⁹

- | | | |
|------|--|--------------------|
| | | $f'_e = [f_e + A]$ |
| (17) | If the miners are in A (and we know it), ... | |
| | a. ... we ought to block A. | ✓ |
| | b. ... we ought to block neither. | ✗ |

This is an *informed* modal base in the sense of §2.3.2 which, as we showed above, yields the same predictions about the *ought*-sentences as the circumstantial one. Thus

²⁹This is arguably the most straightforward implementation of Kratzer's idea: The modal base that gets passed down in the recursive interpretation of the conditional is the result of adding the antecedent to the original modal base. But this view may be too simplistic. The context generally provides more than one epistemic modal base (see Footnote 22 for references), and a lot depends on the details of the selection among them. One often-made proposal for the CCO construal is that a mechanism akin to anaphora resolution makes the restriction of the outer modal to the antecedent-worlds carry over to the inner modal. Frank (1996) and Geurts (1999) offer independent motivation for the anaphoric nature of modal bases. Owing to these complexities, we refrain from claiming that (17) represents Kratzer's own analysis.

the predictions here are analogous to those above.

In contrast, in (18) the modal base for *ought* is the uninformed modal base f_e relative to which the entire conditional is interpreted.

- | | | |
|------|--|--------------|
| (18) | If the miners are in A (and we don't know it), ... | $f'_e = f_e$ |
| | a. ... we ought to block A. | X |
| | b. ... we ought to block neither. | X |

In principle, we must distinguish $f_e(w_A)$, the value of this modal base at an arbitrary antecedent-world w_A at which the consequent is evaluated, from $f_e(w)$, the value of the same modal base at the world at which the conditional is evaluated. However, under the very reasonable assumption that f_e is *introspective*,³⁰ the respective modal backgrounds are the same: $\bigcap f_e(w_A) = \bigcap f_e(w)$, hence also $\bigcap f'_e(w_A) = \bigcap f_e(w)$. As a result, we obtain the same predictions as with the uninformed epistemic reading of unembedded *ought* discussed in §2.3.2.

These predictions take us part of the way towards an account of the non-reflecting reading of the conditionals (which does not figure in Kolodny and MacFarlane's discussion of the miners puzzle): We correctly predict the falsity of (18-a). However, we cannot account for the intuitive truth of (18-b). Noting that the modal background and ordering source for *ought* are the same in the conditional in (18-b) as in the standalone sentence in (2-a), we see that we have run up against the same problem as before: Blocking neither shaft does not come out as the preferred action as long as there are worlds in the modal background at which we save all miners by blocking a shaft.

2.3.4 Interim summary

Let us take stock. Holding fixed the ordering source given in (10), we considered a number of modal base choices for the interpretation of *ought* in non-conditional and conditional sentences.

We easily derived the objective reading by evaluating *ought* relative to a circumstantial modal base (cf. §2.3.1). However, we encountered two problems with the deliberative reading of standalone *ought* and with conditionals: First, interpreting non-conditional sentences relative to an uninformed epistemic modal base, we predicted that all of the following sentences are false in the scenario³¹ (cf. §2.3.2), but this prediction is wrong for (2-a).

- (19) We ought to block shaft A (B).
 (12) There is a shaft that we ought to block.
 (2-a) We ought to block neither shaft.

Second, in conditionals, while we can account for the reflecting reading on

³⁰A modal base f is introspective iff for all worlds w, v , if $v \in \bigcap f(w)$ then $\bigcap f(v) = \bigcap f(w)$. See Fagin et al. (1995) for a basic introduction and Kaufmann et al. (2006); von Fintel and Gillies (2010); Gillies (2010) for arguments that introspection is a reasonable condition to impose on epistemic (or doxastic) modal bases in general.

³¹Recall that (12) is true on the objective reading.

Table 4: Predicted truth values relative to three modal bases at A-worlds.

<i>Modal Base for ought:</i>	<i>uninformed epist.</i>	<i>informed epist.</i>	<i>circumstantial</i>
we ought to block neither shaft	✗	✗	✗
we ought to block shaft A	✗	✓	✓

which (2-b) and (2-c) are true, none of the parameter settings we considered delivers the non-reflecting reading on which (5) is true (cf. §2.3.3).³²

- (2-b) If the miners are in shaft A, we ought to block shaft A.
- (2-c) If the miners are in shaft B, we ought to block shaft B.
- (5) If the miners are in shaft A (B), we (still) ought to block neither shaft.

The crucial facts are collected in Table 4. As before, the symbols ‘✓/✗’ stand for ‘predicted to be true/false under the theory’. The table does not distinguish between non-conditional and conditional contexts, since the structural environment does not determine the modal base for *ought*: We discussed all three possibilities in both unembedded and embedded contexts.³³

The boxed cell in the upper left corner corresponds to both the uninformed deliberative reading of *ought* in standalone sentences and its non-reflecting deliberative reading in conditionals. This is the cell – the only one, we might point out – in which the predictions of the Kratzer-style account are problematic. It is problematic in both non-conditionals like (2-a) and conditionals like (5). Indeed, it is now obvious but nonetheless worth emphasizing that the problem, from the perspective of our analysis, is first and foremost about different modal bases – uninformed *vs.* informed – for *ought*. While superficially there seems to be a correlation with the distinction between non-conditionals and conditionals, respectively, on closer scrutiny this correlation turns out to be imperfect and not at the heart of the problem. Consequently, our focus in the remainder of the paper will not be on conditionals.

That said, before moving on we would like to point out that close reflection on the correlation the uninformed/informed distinction on the one hand and the standalone/conditional distinction on the other unearths a number of disconcertingly important unresolved questions about the interpretation of conditionals.

The general Kratzerian idea that *if*-clauses restrict modal backgrounds remains unquestioned; but how exactly does this restriction come about? Most authors assume that the addition of the antecedent to the modal base simulates, as it were, the

³²Recall that as far as the predictions about *ought*-sentences are concerned, circumstantial and informed epistemic modal bases are not distinguishable. Therefore we actually have no basis for deciding whether the reflecting reading in conditional consequents is deliberative (i.e., informed epistemic) or objective (i.e., circumstantial). We leave this question open. However, we take it to be obvious that the non-reflecting interpretation is deliberative.

³³The informed epistemic reading was relevant for conditionals under the OCO construal as well as the reflecting epistemic variant of the CCO construal. Outside of conditionals, we briefly discussed an informed epistemic reading in §2.3.2.

agents' reaction upon learning that the antecedent is true. However, as we argued in Fn. 23 above, with regard to the miners scenario, this account leaves something to be explained. The judgment that the agents ought to block neither shaft crucially relies on the assumption that finding out is not an option. Does this mean that they believe that it is *impossible* to find out (and thus the update would require a retraction to maintain consistency)? Or can we say that they consider it possible to find out, but don't know how to get the information? In the latter case, the possibility of their finding out by serendipity does not pose a problem for the theory; but it still raises the question how to characterize the agents' inability to attain the information by design, which seems so central to the judgment that (2-a) is true.³⁴

Another open issue becomes pressing when we consider modals in the consequent. The availability of both a reflecting and a non-reflecting reading for such modals suggests an amount of flexibility and context-dependence that Kolodny and MacFarlane would seem hard pressed to incorporate in their account. It may be considered good news that the full-blown Kratzer-style approach is rich enough to accommodate these distinctions; but one man's increase in versatility and expressiveness is another man's proliferation of spurious ambiguity. The question of how exactly the parameters for the interpretation of modals in conditional consequents are set remains open and continues to require careful empirical investigation and theory-building.

2.4 Introducing information dependence

Kolodny and MacFarlane's (2010) treatment of modal verbs is similar to Kratzer's in many respects, but deliberative *ought* in the miners case poses no problem for it. In this section we briefly compare the two accounts in order to bring out the crucial difference between the two. By way of preview, this difference consists in *serious information dependence*, a property that Kolodny and MacFarlane's (2010) account has and Kratzer's lacks (or so we argue). Nevertheless, as we point out in Section 2.4.3, Kratzer's framework has some advantages which we do not want to sacrifice. For this reason, we ultimately develop a variant of the Kratzer-style framework that embraces serious information dependence.

2.4.1 Information dependent *ought*

Like Kratzer, Kolodny and MacFarlane interpret modals by quantification over a set of possible worlds whose identity is subject to two contextual parameters.³⁵ The first is an *information state*, formally represented as a set i of worlds. The second is a *selection function* mapping i to a set of worlds which then constitutes the domain of quantification.

Both epistemic and deontic modals are sensitive to the information state, but they differ in the selection function they invoke. An *epistemic* selection function e triv-

³⁴It is interesting in this connection that in some languages like German, the analog of (2-b) would require subjunctive mood. We leave a detailed investigation of this phenomenon for future work.

³⁵They describe this feature as an important respect in which their approach differs from "the usual approach." By the latter they evidently do not mean the prevalent approach in the linguistic literature.

ially selects the information state itself: $e(i) = i$ for all i . A *deontic* selection function d is more flexible, subject only to a constraint which Kolodny and MacFarlane call “realism” and spell out as the condition that $d(i) \subseteq i$ for all i . The intention here is that only worlds are selected that might, as far as the information in i is concerned, be the actual one.³⁶

Thus far there appears to be no semantically significant difference between Kratzer’s framework and Kolodny and MacFarlane’s. Conceptually, the modal background of the former corresponds to the information state of the latter, while Kratzer’s ordering source corresponds to Kolodny and MacFarlane’s selection function: Both identify the domain of worlds over which the truth definitions for the modal operators quantify. The basic form of the truth conditions is quite similar. Restricting our attention for the sake of simplicity to modal sentences $\Box\varphi$ whose prejacent φ does not itself contain any modals, we can state the truth conditions as follows:³⁷

- (20) a. $\Box_s\varphi$ is true at $\langle i, w \rangle$ iff for all $v \in s(i)$, φ is true at $\langle i, v \rangle$.
 b. $\Box\varphi$ is true at $\langle f, g, w \rangle$ iff for all $v \in O(f, g, w)$, φ is true at $\langle f, g, v \rangle$.

Nor does there seem to be a significant difference in the interpretation of conditionals. Kolodny and MacFarlane’s definition (21) is similar to the OCO construal in the Kratzer-style framework.³⁸ Both call for an evaluation of the modalized matrix clause relative to the set of antecedent-worlds in the original information state (21) or modal background (14).

- (21) $[\text{IF}\varphi]\Box_s\psi$ is true at $\langle i, w \rangle$ iff $\Box_s\psi$ is true at $\langle i', w \rangle$,
 iff for all $v \in s(i')$, ψ is true at $\langle i', v \rangle$,
 where $i' = \{w' \in i \mid \varphi \text{ is true at } \langle i, w' \rangle\}$

- (14) $[\text{IF}\varphi]\Box\psi$ is true at $\langle f, g, w \rangle$ iff $\Box\psi$ is true at $\langle [f+\varphi], g, w \rangle$,
 iff for all $v \in O([f+\varphi], g, w)$, ψ is true at $\langle [f+\varphi], g, v \rangle$

Moreover, for both epistemic and deontic modal operators, both frameworks are set up so as to guarantee that the set of worlds over which the quantification ranges is a subset of the modal background (by design, since $O(f, g, w) \subseteq \bigcap f(w)$ for any g) or the information state (by stipulation, since both $d(i)$ and $e(i)$ are subsets of i). Clearly, then, the reason why Kratzer’s framework but not Kolodny and MacFarlane’s runs into problems with deliberative *ought* must lie in the way in which the relevant subset is identified.

³⁶Note that this is different from the notion of “realism” introduced in Kratzer (1981) which in Kolodny and MacFarlane’s framework would amount to the requirement that i contain the world of evaluation.

³⁷The subscript ‘ s ’ on the modal operator in (20-a) is a variable ranging over selection functions. In Kolodny and MacFarlane’s notation it serves the dual role of marking the modal operator as sensitive to a selection function and of indicating which selection function it is. In our notation, the first is achieved by the ‘H’ inside the operator and the second by the ‘ g ’ in the tuple of parameters of evaluation. The difference is merely notational.

³⁸We give a simplified variant of Kolodny and MacFarlane’s truth conditions (their “first approximation”) which they endorse only for antecedents that do not contain epistemic modals. For the general case, they adopt a slightly more complex version inspired by Yalcin (2007).

2.4.2 Serious information dependence

In both frameworks, ought is interpreted as a quantifier whose domain is selected from a set of worlds determined by some body of information. The crucial difference lies in the dependencies between the parameters driving this selection. In Kratzer's account, once the ordering source g is fixed by the context, its value depends only on the world w of evaluation: the induced ordering \leq is unaffected by the value of the modal base. This has an important consequence: Consider two arbitrary modal bases f_1, f_2 such that the corresponding modal backgrounds $\bigcap f_1(w), \bigcap f_2(w)$ both contain two worlds u, v . It is impossible for $O(f_1, g, w)$ to contain u but not v while $O(f_2, g, w)$ contains v but not u , for the former implies that $u <_{g(w)} v$ while the latter implies that $v <_{g(w)} u$, which cannot both hold.

Kolodny and MacFarlane's account is not subject to this constraint, since here the selection of best worlds from the information state is not grounded in independent criteria. For two information states i_1, i_2 , both of which contain the worlds u, v , it is perfectly possible for $d(i_1)$ to contain u but not v while $d(i_2)$ contains v but not u .

As evidenced by the miners scenario, this flexibility is needed in the analysis of deliberative modality: The relative goodness of the possible actions – blocking A, B, or neither – can be reversed depending on how much information is available. To allow for such reversals in relative goodness, Kolodny and MacFarlane require deontic selection functions to have the following property.

Definition 1 (Serious information dependence – Kolodny and MacFarlane)

A deontic selection function d is seriously information-dependent iff for some i_1, i_2 such that $i_2 \subseteq i_1$, there is a world $w \in i_2$ such that $w \in d(i_1)$ but $w \notin d(i_2)$.

To show that Kratzer's account lacks this property, we need to make our informal talk a bit more precise. So far we were putting various ingredients of the two frameworks side by side, but there is strictly speaking no exact correspondence between them: Kolodny and MacFarlane's i and Kratzer's f are of different types, as are the selection function d and the ordering source g . We can, however, define an auxiliary notion of **Kratzer selection function** in terms of g and w that is exactly parallel to Kolodny and MacFarlane's. This is possible because, once the value of the ordering source is fixed, the set of best worlds is determined by the modal background:

Fact 1

For all possible worlds w and modal bases f, f' : If $\bigcap f(w) = \bigcap f'(w)$, then for all ordering sources g , $O(f, g, w) = O(f', g, w)$. Thus in particular, for all f, g, w , $O(f, g, w) = O(\lambda v. \{\bigcap f(w)\}, g, w)$.

This follows from (7) above, where $O(f, g, w)$ was defined solely in terms of \leq and $\bigcap f(w)$. We can now define, for any g and w , a function mapping modal backgrounds to the corresponding sets of best worlds under $g(w)$.

Definition 2 (Kratzer selection function)

For all worlds w and ordering sources g , the Kratzer selection function $d_{g,w}$ is a function from propositions to propositions such that for all $p \subseteq W$, $d_{g,w}(p) = O(\lambda v. \{p\}, g, w)$.

Theorem 1

Table 5: Selection function in the miners scenario

<i>Information state</i>	<i>Best Worlds</i>
$i = AA \cup AB \cup AN \cup BA \cup BB \cup BN$	$d(i) \subseteq AN \cup BN$
$i_A = AA \cup AB \cup AN$	$d(i_A) \subseteq AA$
$i_B = BA \cup BB \cup BN$	$d(i_B) \subseteq BB$

There is no ordering source g and possible world w such that the Kratzer selection function $d_{g,w}$ is seriously information-dependent.

Proof. Suppose for reductio that $d_{g,w}$ is seriously information-dependent for some g, w . Thus there are information states i_1, i_2 such that $i_2 \subseteq i_1$ and for some world $v \in i_2$, (i) $v \in d_{g,w}(i_1)$ but (ii) $v \notin d_{g,w}(i_2)$. By (ii), there is a world $u \in i_2$ such that $u < v$. But $u \in i_1$ since $i_2 \subseteq i_1$, hence $v \notin d_{g,w}(i_1)$, contradicting (i). \square

Theorem 1 brings out the crux of the problem with Kratzer-style semantics and deliberative *ought*. We cannot find values for the modal base f_e and the ordering source such that, at the world of evaluation w , blocking neither shaft is the best course of action relative to those parameters while blocking a shaft is best relative to a more informed modal base f'_e such that $\bigcap f'_e(w) \subseteq \bigcap f_e(w)$. Kolodny and MacFarlane's ability to account for the fact that all of (2-a), (2-b) and (2-c) are true rests crucially on the ability to make this switch.

2.4.3 In favor of transparency

Kolodny and MacFarlane clear the path towards serious information dependence, and hence to a satisfactory analysis of deliberative *ought*, by rejecting the independence of the information state and the criteria for prioritizing worlds. Kratzer's ordering source ranks worlds once and for all (relative to a given world of evaluation), thus while different modal bases may yield different sets of best worlds, those best worlds are in all cases determined by the same ordering relation. In contrast, Kolodny and MacFarlane's key idea is that how worlds compare with each other can depend on the information state. However, they make no attempt to *derive* the set of best worlds from underlying preferences that are constant and independent of the information state.

In the miners scenario, the intended behavior of the selection function can be diagrammed as in Table 5. As there is no independent information about what the selection function should look like, d is assumed to satisfy the desired predictions by stipulation. We agree that this is how the best worlds should be selected in the scenario, and we saw above that this poses a challenge for the standard Kratzerian framework. Yet we do not draw the conclusion that we should do away with the ordering source altogether.

Abandoning ordering sources in favor of a more unconstrained selection function would introduce a measure of opacity in the semantic account. As a consequence,

some systematic explanations that were available in the standard framework would be lost. Consider this pair:

- (22) If you are healthy, you ought to read that article.
- (23) If you have high fever, you ought not to read that article.

On the standard Kratzer account, the joint acceptability of these two sentences can be explained systematically, in terms a single initial pair of modal base and ordering source, and their interaction under different restrictions on the former. Let us imagine that at the world of evaluation, the ordering source contains the propositions (i) {that you acquire the information in the article}, and (ii) {that you do not needlessly tire your eyes}; while the modal base specifies that high fever prevents one from acquiring any information. Then if in evaluating (22) we restrict the modal background to the worlds at which you are healthy, reading doesn't tire your eyes needlessly and you can (and therefore should) acquire the information in the article. In interpreting (23), in contrast, the modal background is restricted to worlds in which you have high fever. None of those worlds satisfy the proposition that you acquire the information in the article, and moreover at all of them reading tires your eyes needlessly. Hence the reading-worlds are ranked below the non-reading worlds.³⁹

The account we get from Kolodny and MacFarlane is not nearly as systematic as that. The workings of the selection function remain inaccessible. We take it that it would be preferable to have a systematic account of the miners case, similar to the one we just sketched for the joint truth of (22) and (23).

3 Extending the standard analysis

In §2.1 we showed that the standard Kratzer-style translation of *ought* as a human necessity operator with an ordering source reflecting the agents' preferences fails to predict the truth of (2-a), regardless of whether the modal base is circumstantial or epistemic. The reason is that worlds at which the agents block neither shaft are invariably outranked by worlds at which they block the shaft in which the miners are. As long as the modal background contains worlds of the latter type, (2-a) cannot be true. Intuitively blocking the shaft in which the miners are is indeed the best possible outcome. But the agents cannot rely on this criterion in choosing among the alternatives available to them because they lack the information which shaft is the one to block: Considering the worlds that are compatible with the description of the scenario, among the ones at which they jump to action and block shaft A there are excellent ones (where all ten miners are saved) and dreadful ones (where all miners drown). Similarly for blocking shaft B. Since finding out where the miners are is not an option, neither of these two actions is "safe," and the third choice of blocking neither shaft (where nine of the miners are saved) outranks both.

While the agents' ignorance is a crucial factor in this reasoning, it is not sufficient to explain the truth of (2-a) (appealing to an uninformed epistemic modal base did not

³⁹Note that in this case the choice between the OCO and CCO construals, as well as various values for the modal bases, does not lead to any interesting differences.

solve the problem, see §2.3.2 above). What is needed in addition to a representation of their limited information is an account of the actions available to them. Intuitively, blocking shaft A and blocking shaft B are among the choices they have, but blocking the shaft in which the miners are is not. Thus while the agents have a clear preference among the possible outcomes, none of the actions available to them is a sure way to secure the best result. To capture the deliberative reading of sentences like (2-a), our semantic interpretation must be made sensitive to constraints of this sort.

3.1 Decision problems

To this end, we propose an extension of the standard Kratzerian framework. In addition to the usual two parameters (modal base and ordering source), deliberative modality is sensitive to a contextually salient *decision problem* δ , identifying at each world the set of actions from which the agent has to choose at that world.⁴⁰ For simplicity, we assume that in our possible-worlds framework each action can be represented as the set of worlds in which that action is taken – i.e., formally, actions are propositions.⁴¹

Decision problems thus understood need not partition the logical space, or even the set of epistemically possible worlds. For one thing, the alternatives need not be logically incompatible. There are worlds at which we block both shaft A and shaft B, but those are ruled out by the description of the scenario. The actions are only mutually incompatible given the situation; hence we require that the corresponding propositions be mutually exclusive relative to the modal background.^{42,43} We simply assume that while the contextually decision problem need not partition the logical space, it has to determine a partition of the modal background relative to which the

⁴⁰Note that we conceive of the decision problem as world dependent. Technically, this makes it just another conversational background, i.e., a function from worlds to sets of propositions. We take this move to be warranted in view of examples like (i):

- (i) If blocking shaft A and B were an option, we ought to do that.

For a non-counterfactual example, consider a variant of the miners problem in which the agents are unsure as to whether they have a sufficient number of sandbags to block both shafts simultaneously.

- (ii) If blocking both shafts is an option, we ought to do that.

The indicative variant in (ii) appears to be true in the scenario thus modified.

⁴¹Based on a very different motivation, Cariani (2009, forthcoming) discusses a formal framework that makes use of decision problems, but diverges from ours on key resources and motivational points. The idea of representing actions as propositions and decision problems as sets of proposition has a long history, but was recently articulated in the context of the logic of action and deontic logic by Horty and Belnap (1995); Belnap et al. (2001); Horty (2001).

⁴²Nor do the actions constituting a decision problem have to jointly exhaust all possibilities, even those that are epistemically accessible. At some possible worlds, the agent does not act at all (she may be struck by lightning or a meteor, for example). There are various ways of dealing with this possibility (either semantically or pragmatically), but adjudicating the issue is largely orthogonal to our aims here.

⁴³The formal similarity between decision problems in our model and question denotations is evident. The possibility for the former to leave out a non-empty set of worlds at which no action is taken is reminiscent of the *residual answer* in the latter (Hamblin, 1958). It is customary in question semantics to set the residual answer aside by assuming that questions in actual usage come with an implicit presupposition-like domain restriction.

deliberative modal is interpreted.⁴⁴ Formally, we use ‘ $[\delta|f]$ ’ to refer to a conversational background which results from combining the modal base f with the decision problem δ . At each world, $[\delta|f]$ returns a restriction of the decision problem to the worlds in the modal background:

Definition 3 (Filtered decision problem)

Let δ be a decision problem and f a modal base. The result of filtering δ by f is a decision problem $[\delta|f]$ defined as follows: For all worlds w ,

$$[\delta|f](w) = \{\bigcap(f(w) \cup \{p\}) \mid p \in \delta(w)\} \setminus \emptyset$$

It is easy to show that if the propositions in $\delta(w)$ are mutually exclusive, then so are the propositions in $[\delta|f](w)$, for any f .

For deliberative *ought*, the decision problem spells out which propositions are *choosable* for the agent.⁴⁵ In particular, in the miners scenario, the natural decision problem is (24):

$$(24) \quad \delta(w) = \{\text{we block shaft A, we block shaft B, we block neither shaft}\}$$

Since we take δ to be a contextual parameter alongside f and g , we now represent contexts as triples of the form $\langle f, g, \delta \rangle$. As we anticipated, to simplify the notation, we occasionally use the variable ‘ c ’ to range over such triples, writing ‘ f_c ’ etc. to refer to the corresponding components. As before, in addition to contexts we have points of evaluation (used in the recursive truth definitions) which are represented by adding a world parameter to a context – that is to say, if $\langle f, g, \delta \rangle$ is our context, a point of evaluation is a quadruple of the form $\langle f, g, \delta, w \rangle$.

To take into account the information dependence of deliberative modality, the modal base is epistemic. The ordering source fixes the preferences underlying deliberation, as before. What is novel about our approach is not our choice of these two parameters, but the way in which an ordering relation on possible worlds is derived from them: We coarsen the order so that it does not distinguish between worlds in one and the same cell of the decision problem. In other words, for each cell in the decision problem, all worlds in which the corresponding action is taken are tied for relative goodness or badness. Formally, for a context $c = \langle f, g, \delta \rangle$ and world w we

⁴⁴This is required in the definition of the pre-order on which the modal verb relies (cf. 6). In that sense, the requirement that the contextually salient decision problem partition the contextually salient modal background can be seen as a (semantic) presupposition of the modal verb.

⁴⁵Roughly speaking a choosable proposition is one which the agent can bring about while being aware that she is doing so. Thus Kolodny and MacFarlane (2010) define (p. 20):

Definition 4 (Chooseable (preliminary – Kolodny and MacFarlane))

ϕ is choosable relative to $\langle i, w \rangle$ iff there is some action specification Δ such that ‘ $\diamond_c(\Delta$ is done by agents who know they are doing Δ)’ and ‘ $\square_c(\Delta$ is done $\supset \phi$)’ are both true at $\langle w, i \rangle$.

But this seems to us too weak on the first conjunct: It only requires that the agents hold it possible to do Δ (knowingly). We propose the following alternative (where \diamond_c stands for circumstantial possibility):

Definition 5 (Chooseable)

ϕ is choosable relative to $\langle w, i \rangle$ iff there is some action specification Δ such that ‘ $\square_c(\diamond_c(\Delta$ is done by agents who know they are doing $\Delta))$ ’ and ‘ $\square_c(\Delta$ is done $\supset \phi$)’ are both true at $\langle w, i \rangle$.

The issue deserves further discussion, which we leave for another occasion, however.

define an order $\leq_{c,w}$ which compares worlds only in terms of ordering source propositions that hold throughout their respective cells in the partition induced by $[\delta_c|f_c](w)$.

First two pieces of notation. We refer to the set of worlds at which the same action is taken as at some given world v as follows:

(25)

$$[v]_{c,w} := \begin{cases} \text{the unique } p \in [\delta_c|f_c](w) \text{ s.t. } v \in p & \text{if } v \in \bigcap f_c(w), \\ W & \text{otherwise.} \end{cases}$$

This ensures that the mapping $[\cdot]_{c,w}$ is defined for all worlds, while also preventing worlds outside of the modal background from outranking ones within. Second, we define a function $\blacktriangle_{c,w}$ from worlds to sets of propositions as follows.

$$(26) \quad \blacktriangle_{c,w}(v) = \{p \in g_c(w) \mid [v]_{c,w} \subseteq p\}$$

Intuitively, $\blacktriangle_{c,w}(v)$ picks out the set of ordering source propositions that ‘count in favor of’ v . For a proposition to count in favor of a world v , it is not sufficient that it be true at v . Rather, its truth must be *guaranteed* by the action taken at v . Only those ordering-source propositions for which this is the case can affect v ’s relative goodness. Notice that if v is not in the modal background, then only tautologies count in favor of it, and those, of course, count in favor of all worlds. The resulting ordering is defined as follows:

Definition 6 (Deliberative preference ranking)

For any context c and world w , the deliberative preference ranking is a binary relation $\leq_{c,w}$, defined as follows, for all $u, v \in W$:

$$u \leq_{c,w} v \text{ iff } \blacktriangle_{c,w}(v) \subseteq \blacktriangle_{c,w}(u)$$

As in the case of Kratzer’s $\leq_{g(w)}$, it is easy to see that $\leq_{c,w}$ is a pre-order (i.e. that it is reflexive and transitive).

The key difference between the two is the following: For a proposition $p \in g_c(w)$ to have any effect at all on the ordering of worlds, there must be some cell in the decision problem which (given the agent’s information) will secure p . Furthermore, p only counts in favor of worlds in such cells: The relative standing of worlds in other cells (at which p is not guaranteed) is unaffected even if p is true at them. Thus in effect, in comparing the various choices given by the decision problem, a cell cannot be better than its worst conceivable outcome (in light of the agent’s information).

The truth conditions of deliberative *ought* are analogous to the ones given before:

$$(27) \quad \text{a. } O(c, w) = \{u \in \bigcap f_c(w) \mid \neg \exists v \in \bigcap f_c(w) [v <_{c,w} u]\}$$

$$\text{b. } \Box\varphi \text{ is true at } \langle c, w \rangle \text{ iff for all } v \in O(c, w), \varphi \text{ is true at } \langle c, v \rangle.$$

Notice that as long as $g_c(w)$ is finite, there is guaranteed to be a set of best worlds under (27-a), even if the decision problem is not finite.

Table 6: Parameters for the deliberative interpretation with an uninformed modal base

	w
modal background	$\bigcap f_e(w) = AA \cup AB \cup AN \cup BA \cup BB \cup BN$
ordering source	$g(w)$
decision problem	$\delta = \{\text{block A, block B, block neither}\}$
filtered problem	$[\delta f_e](w) = \{AA \cup BA, AB \cup BB, AN \cup BN\}$
ranking	$AN \cup BN <_{c,w} AA \cup BA, AB \cup BB$
minimal words	$O(c, w) = AN \cup BN$

3.2 Predictions

The predictions for the miners scenario are given in Table 6. Recall that the ordering source is defined as follows:

$$(10) \quad g(w) = \{\text{all miners are saved, at least 9 miners are saved, } \dots, \\ \text{at least 1 miner is saved} \}$$

Under uncertainty, the set of worlds in which we block shaft A contains both worlds at which the miners are in shaft A (and we save all of them) and worlds at which they are in shaft B (and we lose all of them). Although the former are ranked highest under $g(w)$, their good standing does not carry over to the modified ranking once the decision problem is taken into account. This is because none of the three available actions guarantees that the highly desirable outcome actually comes to pass: The equally possible outcome of losing all miners drags down the ranking of the good worlds in those cells in which we block shaft A (similarly for shaft B). In contrast, at all worlds in which no shaft is blocked, nine miners are saved no matter what. As a result, the worlds in this cell strictly outrank those in the other two.

Things would change if we were to find out (somehow, to our surprise – see §2.3.2) where the miners are: Consider for concreteness the case of learning that that they are in A. The resulting parameters are shown in Table 7. The modal background at w is restricted to the A-worlds in $\bigcap f_e(w)$, thus the filtered decision problem becomes $[\delta|[f_e+A]]$ – each of the three cells is restricted to A-worlds. This changes which ordering-source propositions play a role in determining the relative goodness of worlds: Blocking A is now guaranteed to save all ten miners, thus the proposition that all ten are saved counts in favor of those world at which it is true. As a result, the worlds at which we block shaft A now strictly outrank all others. This is what we wanted.

Our analysis shares with Kolodny and MacFarlane’s the feature that the selection of optimal worlds is *seriously information dependent*. However, we maintain that it provides a more specific and systematic diagnosis.

Kolodny and MacFarlane gloss their notion of serious information dependence as the idea that deontic ideality changes as new information becomes available. On

Table 7: Parameters for the deliberative account with an informed modal base

	w
modal background	$\bigcap f_e(w) = AA \cup AB \cup AN$
ordering source	$g(w)$
decision problem	$\delta = \{\text{block A, block B, block neither}\}$
filtered problem	$[\delta f_e](w) = \{AA, AB, AN\}$
ranking	$AA <_{c,w} AN <_{c,w} AB$
minimal words	$O(c, w) = AA$

our view there is a sense in which this is correct and a sense in which it is not. In particular, while it is true that the ranking and the optimal worlds vary with the information at the agents’ disposal, the ordering source g remains the same. The variation in our ranking of AA worlds does not depend on the (broadly speaking) normative component of the model (i.e. the ordering source). It rather depends on the fact that the same action (e.g. blocking shaft A) is associated with different sets of worlds under different information states (for example, in the uncertain state blocking shaft A corresponds to {AA, BA}, whereas if we know that the miners are in shaft A, it is {AA}).

3.3 Relation to Decision Rules.

In principle, our approach is compatible with a number of different ways of comparing actions. Some people [names of omitted for blind review] have questioned whether our semantics encodes the *MaxiMin* principle – the decision theoretic rule that requires agents to evaluate actions in terms of their worst conceivable outcome and choose the “least bad” one among them. This observation, if correct, may be construed as an objection because, the argument goes, the semantics should not encode such decision-theoretic principles (either because the principle is wrong, or because it isn’t the job of the semantics to settle it).

We think this charge is partly wrong and partly right. It is partly wrong because our formal structure is quite different from that of decision theory. In particular, we have not suggested that there is a distribution of *utility* on the individual worlds and generally nothing in our apparatus plays the role of outcomes in decision theory. Finally, for these reasons, we have not advanced a principle requiring those choices that maximize the minimum outcome. The objection is partly right, however, because like MaxiMin, our recipe for comparing actions is only sensitive to those priorities that can be guaranteed.

Faced with this issue, one might simply remark that in the same sense in which our semantics encodes MaxiMin, the classical semantics encodes *MaxiMax*. Alternatively, one might consider an approach in terms of expected utilities. We do think such an approach is promising, and we do think that it is an accessible generalization of our current framework. We refrain, however, from developing it here, as we are

investigating the prospects of an analysis in a purely qualitative framework.

We can, however, inject some flexibility into our approach by revising the definition in (26), which we repeat here for ease of reading:

$$(28) \quad \blacktriangle_{c,w}(v) = \{p \in g_c(w) \mid [v]_{c,w} \subseteq p\}$$

Instead of scoring the worlds in the cell $[v]_{c,w}$ according to which ordering-source propositions the action guarantees, we could require that some other, possibly contextually determined relation R_c must hold between $[v]_{c,w}$ and those propositions.⁴⁶

$$(29) \quad \blacktriangledown_{c,w}(v) = \{p \in g_c(w) \mid R_c([v]_{c,w}, p)\}$$

In our account, R_c is instantiated with the subset relation (i.e., $[v]_{c,w}$ must entail p in order for p to count in favor of the worlds in $[v]_{c,w}$). But other values are possible. For instance, R_c could hold whenever $[v]_{c,w}$ and p are compatible (thus ‘encoding MaxiMax’), or when p is a weak Kratzerian necessity given $[v]_{c,w}$ relative to a stereotypical ordering source that is separate from the deontic one we have been discussing.

We will leave this issue here and defer development of a richer semantics to separate work. In the remainder of the paper we work with the ordering generated by \blacktriangle rather than by the more flexible \blacktriangledown .

3.4 Circumstantial modality and stability

To account for the information dependence of *ought* on its deliberative reading, we appealed to an epistemic modal base. Now that the analysis is in place, it is easy to see that if we were to use a circumstantial modal base instead, we would not predict the relevant judgments, despite the introduction of the decision problem as a third parameter.

The reason is that, as we argued in §2.3.1 above, the value of the circumstantial modal base f_o at any world w includes the proposition that the miners are in the shaft in which they are at w , with the result that they are in the same shaft at all worlds in the modal background. Thus for instance at an A-world w_A , $\bigcap f_o(w_A) = AA \cup AB \cup AN$; hence the filtered decision problem is $[\delta|f_o](w) = \{AA, AB, AN\}$; thus clearly $AA <_{c,w} AN <_{c,w} AB$ (just as with the informed epistemic modal base $[f_e+A]$ above). So the prediction is that we ought to block shaft A, not neither shaft. Analogously, for a world in which the miners are in shaft B it is predicted that we ought to block shaft B. In the given scenario, these predictions could not be reversed by adding further information to the modal base consistently.

We think that this stability should follow from general assumptions about the circumstantial modal base. As we mentioned above, there is a certain vagueness in the standard Kratzerian definition of the circumstantial conversational background as assigning to each world a specification of the “relevant circumstances” at that world. But while it may not always be clear what exactly the circumstantial modal base looks like, it does seem reasonable to expect that it should at least contain all the in-

⁴⁶Carr (ms.), who raises the MaxiMin objection, defends a similar view in a probabilistic setting.

formation needed to resolve the contextually given problem (if there is one) in w . To begin to formalize this notion, we introduce a property which we call *persistence*.⁴⁷

Definition 7 (Persistence)

Let $c = \langle f, g, \delta \rangle$. Then f is persistent for δ and g at a world w if and only if for any proposition $p \subseteq W$ and worlds $u, v \in \cap [f+p](w)$,

$$u \leq_{[c+p],w} v \text{ iff } u \leq_{c,w} v$$

One way in which persistence arises is for all propositions in the ordering source to count in favor of all worlds in the modal background at which they are true. This is the case whenever none of these propositions “cuts across” any cells of the filtered decision problem. We state this fact without proof.⁴⁸

Fact 2

A modal base f is persistent for δ and g at w if for all $v \in \cap f(w)$ and $p \in g(w)$, if $[v]_{c,w} \cap p \neq \emptyset$ then $[v]_{c,w} \subseteq p$.

The notion of persistence is of course equally relevant to an epistemic modal base: It specifies whether or at what point agents have acquired enough information to resolve their decision problem under full consideration of their criteria. In particular, when an epistemic modal base is persistent for δ and g at w , this means that the judgments for deliberative *ought* with respect to these parameters coincide with the ones for two-parameter *ought* with the same g and a circumstantial modal base, which we considered the appropriate rendering for objective *ought*.

4 Relation to the standard Kratzer analysis

The standard Kratzerian framework relies on two parameters, the modal base and the ordering source. We added a decision problem as a third parameter and obtained more accurate predictions about deliberative *ought* than the two-parameter analysis we used for comparison. But is the third parameter really necessary, and just how big a departure from the standard picture is its introduction? These are important questions, but they don’t receive clear answers based on linguistic data alone. Rather, our response appeals to certain conceptual and theoretical preferences and concerns which we consider sensible, though not without alternative.

In this section we reflect on some of the choices involved: First, does the introduction of a third parameter for deliberative *ought* commit us to a lexical ambiguity between deliberative (three-parameter) and non-deliberative (two-parameter) *ought*? Second, does appealing to a specific parameter in evaluating deliberative *ought* undermine claims about the unity and interrelation of the modal system? Third, is a two-parameter analysis of deliberative modality altogether impossible or just dispreferred, and if the latter, on what grounds? In discussing these issues, we hope to

⁴⁷In Definition 7, ‘ $[c+p]$ ’ is shorthand for $\langle [f+p], g, \delta \rangle$.

⁴⁸The converse does not hold: It is possible for persistence to hold even when some ordering-source propositions are not “operative.” For instance, this happens when the filtered decision problem has only one cell, or when the ordering source contains only singleton propositions.

make a convincing argument for our approach; but at the very least, the discussion should clarify some of the concerns and tradeoffs involved.

4.1 Adding the third parameter

Not all occurrences of *ought* are deliberative. In the context of the miners scenario we discussed an objective reading which gave rise to intuitions about our sentences that differed from those under the deliberative reading. Some authors (e.g., Schroeder, 2010) have defended *ambiguity* theories which draw a lexical distinction between the deliberative *ought* at play in the miners scenario and its non-deliberative counterpart, illustrated in (30-a)-(30-b).

- (30) a. The car ought to start in less than five minutes.
 b. Bill ought to receive the prize.

We want to know whether the third parameter that we postulated for deliberative *ought* might have unwelcome repercussions for the analysis of non-deliberative *ought*, for *any* manifestation of the latter. If there were such interference, we would have to shield non-deliberative *ought* from the third parameter, which would be a step towards an ambiguity analysis.

To be sure, the question whether a decision problem plays any role in the interpretation of (30-a) is beside the point. After all, cars don't make decisions. But we can and should ask more abstractly whether the *coarsening* of the logical space induced by the third parameter affects the analysis of *ought* under this interpretation. It turns out that this question need not worry us. If the decision problem is entirely irrelevant, we can neutralize the parameter by setting it to a default value on which our semantics collapses into Kratzer's, regardless of how the modal base and the ordering source are chosen. We call this special value the *maximally specific* decision problem.⁴⁹

Definition 8 (Maximally specific decision problem)

The maximally specific *decision problem* is the partition $\delta^* = \{\{w\} \mid w \in W\}$.

Now we can show that for any modal base f , ordering source g and world w , the ordering relation among the worlds in the modal background that is obtained according to our rules relative to $\langle f, g, \delta^*, w \rangle$ coincides with that induced under the standard Kratzer-style recipe relative to $\langle f, g, w \rangle$.⁵⁰ To improve readability, let $c^* = \langle f, g, \delta^* \rangle$. The claim then is that $\leq_{g(w)}$ matches $\leq_{c^*, w}$ on the worlds in the modal background.

Theorem 2

For any modal base f and ordering source g , let $c^* = \langle f, g, \delta^* \rangle$. Then for all worlds w, u, v , if $u, v \in \bigcap f(w)$ then $u \leq_{g(w)} v$ iff $u \leq_{c^*, w} v$.

⁴⁹Note that this trivial decision problem must be exempt from the requirement of "choosability" discussed in Footnote 45. We assume that only "choosable" decision problems can be contextually salient with deliberative modals.

⁵⁰The two relations may differ on worlds outside of the modal background, but those worlds do not affect the interpretation of modals that we are interested in.

Proof.

$$\begin{aligned}
u \leq_{g(w)} v &\text{ iff } \{p \in g(w) \mid v \in p\} \subseteq \{p \in g(w) \mid u \in p\} && \text{(Def. (6))} \\
&\text{ iff } \{p \in g(w) \mid \{v\} \subseteq p\} \subseteq \{p \in g(w) \mid \{u\} \subseteq p\} && \text{(set theory)} \\
&\text{ iff } \{p \in g(w) \mid [v]_{c^*,w} \subseteq p\} \subseteq \{p \in g(w) \mid [u]_{c^*,w} \subseteq p\} && \text{(Def. 8)} \\
&\text{ iff } u \leq_{c^*,w} v && \text{(Def. 6)}
\end{aligned}$$

□

The significance of this result is twofold: First, it means that we need not take deliberative *ought* to have its own separate lexical entry, even if it turns out that the interpretation of non-deliberative *ought* does not depend on the decision problem. We can simply assume that, in those other cases, the decision problem is set to δ^* by default.⁵¹ Second, in terms of the formal properties of the framework, adding the third parameter does not subtract from the expressive power of the formalism: Whatever can be done with two parameters can also be done with three.

In principle, nothing prevents other modal operators from also being sensitive to decision problems, assuming that the parameter is fixed lexically to δ^* for modals which do not possess a distinct deliberative reading. This would constitute just another lexical restriction on the values of conversational backgrounds that a particular lexical item is compatible with. Kratzer (1981) presents ample evidence that most modal elements at least in English and German carry some such restrictions. We are thus confident that our account does not jeopardize the status of any claims about the underlying unity of the modal system.

4.2 Removing the third parameter

The previous subsection showed that Kratzer’s two-parameter semantics is a special case of our three-parameter semantics, obtained by setting the decision problem to δ^* . Shifting the perspective, another question suggests itself: Can we derive the predictions of our three-parameter semantics in the two-parameter framework?

Our three parameters jointly serve to induce a pre-order on the possible worlds. So does the ordering source in Kratzer’s system, all by itself. Are there pre-orders that could only be induced with three parameters, not with two? The answer is negative: For any f, g, δ and w , there is an ordering source g' such that the pre-order induced by $g'(w)$ under Kratzer’s definition is the same as that induced at $\langle f, g, \delta, w \rangle$ under our rules. Moreover, this g' can be derived from g (with reference to the other parameters) via a straightforward transformation:

Definition 9 (Trimmed context)

The trimming operation \cdot^\downarrow on contexts is defined as follows, for all $c = \langle f, g, \delta \rangle$:

- a. Let $w \in W$.

⁵¹Recall that we made a similar move above, in §2.1, treating the necessity operator \square as a special case of the human necessity operator \boxplus , obtained by setting the ordering source to the trivial conversational background $\lambda v.\emptyset$. Its use as an ordering source or as a modal base is common practice in the Kratzer-style framework (see also Kratzer 2012 for applications in the analysis of conditionals).

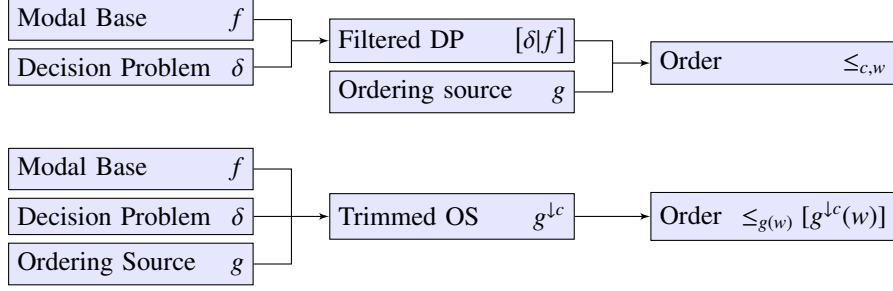


Figure 1: Architecture of our account of deliberative modality (top) and the alternative using a trimmed ordering source (bottom).

- i. For any proposition p , $p^{\downarrow c,w} := \{v \in W \mid [v]_{c,w} \subseteq p\}$.
- ii. For any set A of propositions, $A^{\downarrow c,w} := \{p^{\downarrow c,w} \mid p \in A\}$
- b. $g^{\downarrow c} := \lambda w [g(w)^{\downarrow c,w}]$
- c. $c^{\downarrow} := \langle f, g^{\downarrow c} \rangle$.

The trimmed context c^{\downarrow} is a Kratzer-style, two-parameter context. Its ordering source can be described as “our preferences in view of the information we have about our choices.” Informally, at each world w it is obtained by eliminating from each proposition p in $g(w)$ worlds that are in p but whose cell in the filtered decision problem is not contained in p . Now, instead of defining the ordering between possible worlds in terms of the interplay between g and $[\delta|f](w)$ as we have done so far, we can derive it directly from the trimmed ordering source along the lines of the standard two-parameter Kratzer semantics. The diagrams in Figure 1 contrast the structure of our approach with the structure of the more orthodox alternative just outlined.

It turns out that the two approaches are equivalent for the evaluation of any modal operator that is not modified by a conditional antecedent. That is, the two-parameter approach with c^{\downarrow} yields exactly the same predictions as the three-parameter approach with c .

Theorem 3

For any worlds w, u, v and context $c = \langle f, g, \delta \rangle$, if $u, v \in \bigcap f(w)$ then $u \leq_{c,w} v$ iff $u \leq_{c^{\downarrow},w} v$.

Proof.

$$u \leq_{c,w} v \text{ iff } \blacktriangle_{c,w}(v) \subseteq \blacktriangle_{c,w}(u) \quad (\text{Def. 6})$$

$$\text{iff } \{p \in g_c(w) \mid [v]_{c,w} \subseteq p\} \subseteq \{p \in g_c(w) \mid [u]_{c,w} \subseteq p\} \quad (26)$$

$$\text{iff } \{p \in g(w) \mid v \in p^{\downarrow[\delta|f](w)}\} \subseteq \{p \in g(w) \mid u \in p^{\downarrow[\delta|f](w)}\} \quad (\text{Def. 9a})$$

$$\text{iff } \{p^{\downarrow[\delta|f](w)} \mid p \in g(w), v \in p^{\downarrow[\delta|f](w)}\} \subseteq \{p^{\downarrow[\delta|f](w)} \mid p \in g(w), u \in p^{\downarrow[\delta|f](w)}\}$$

$$\text{iff } \{p \in g(w)^{\downarrow[\delta|f](w)} \mid v \in p\} \subseteq \{p \in g(w)^{\downarrow[\delta|f](w)} \mid u \in p\} \quad (\text{Def. 9b})$$

$$\text{iff } u \leq_{c^{\downarrow},w} v \quad \square$$

The main upshot is that our three-parameter approach does not add to the expressive power of the standard framework of Kratzer-style semantics: Anything that can be done with three parameters (as long as we restrict ourselves to non-conditional *ought*-sentences, see below) can also be done with two. Recall that in Section 4.1 we showed that the converse also holds. Thus formally the two are equivalent.

One could argue now that considerations of conservatism and theoretical parsimony militate against the introduction of a third parameter. However, we believe that there are good reasons to stick with it.

4.3 In favor of three parameters

First, Kratzer’s conversational backgrounds are usually required to be contextually salient (consider in particular implementations in DRT where the conversational backgrounds are explicitly treated as discourse anaphors, e.g. Frank, 1996; Geurts, 1999). The trimmed version considered for deliberative modality does not strike us as a natural object to be contextually salient.

Second, and more seriously, we argued that one of the main advantages of our framework over Kolodny and MacFarlane’s is that the priorities (goals, desires, values etc.) underlying the agents’ deliberation are represented transparently and independently of the information available to them. By forcing our three-parameter deliberative *ought* into the two-parameter mold, we lose this advantage: For instance, learning where the miners are effects a change not only in the modal base, but also in the ordering source. In contrast, if the decision problem is introduced separately, the priorities can be kept invariant, as we think they should be.

Third, the version with the trimmed ordering source runs into a technical problem when it comes to the interaction with conditional antecedents. Remember that in the Kratzer-framework, conditionals can be treated by adding the antecedent proposition to the modal base of the top-level modal in the consequent (under both the OCO and the reflecting CCO construals). In our three-parameter version, as the filtering of the decision problem depends on the modal base, an update of the modal base may affect the order $\leq_{c,w}$ even as $g(w)$ and $\delta(w)$ stay constant. In contrast, in the two-parameter version, for a change in the modal base to affect the ordering, we have to assume that the conditional construction itself shifts the ordering source parameter (in addition to affecting the modal base).⁵²

Another way of seeing the point is this: our proposal and the trimming analysis are equivalent for *unembedded ought*-sentences. However, when *ought* occurs embedded in the consequent of a conditional, and the conditional is given an OCO or reflecting CCO construal and a standard Kratzer-style analysis, the two analyses disagree. To realign them, one would have to postulate that trimming applies in the course of the semantic evaluation (i.e. we first update the modal base with the proposition expressed by the conditional antecedent, and then trim the ordering source, then evaluate the consequent). In our opinion this would constitute a more dramatic departure from the spirit of Kratzer’s approach than our addition of a third parameter.

⁵²Sæbø (2002) develops a theory of anankastic conditionals on which part of the antecedent serves to modify the ordering source of a modal in the consequent. For critical discussion, see von Stechow and Iatridou (ms.).

We conclude that while it is an interesting result that the three-parameter version of the Kratzer-style semantics can be forced into the standard two-parameter framework by using a particular conversational background as the ordering source, we think that there are good arguments to stick with the analysis of deliberative *ought* that explicitly employs a decision problem parameter in addition to modal base and ordering source.

5 Conclusions

The miners problem motivates a seriously information dependent analysis of *ought*. As we have shown, the standard Kratzer analysis cannot be readily applied to it. We have shown that a three-parameter analysis of deliberative *ought* is the key to a satisfactory understanding of serious information-dependence. By adding a decision problem as a third contextual parameter and allowing decision problems to affect the mechanics by which the ordering is generated, we can derive all the salient verdicts in the miners' case. Moreover, we handle serious information dependence without sacrificing the systematicity and transparency of Kratzer's framework

Ultimately, we argued, the introduction of a third parameter constitutes a comparatively insignificant departure from the standard view: any two-parameter solution can be encoded as a three-parameter solution with a trivial decision problem, and our three-parameter version can be translated into a standard two-parameter version. Nevertheless, we believe we have presented convincing arguments that the three parameter version provides a more natural and elegant interface with context.

Acknowledgments

We thank Kai von Fintel, Graham Katz, Angelika Kratzer, John MacFarlane, Aynat Rubinstein, Paul Portner, Paolo Santorio, Malte Willer and Thomas Ede Zimmermann for comments on drafts of the present paper. We also thank the Semantics and Philosophy Workshop at the University of Chicago, the Philosophy Colloquium at the University of Wisconsin, Madison and at the University of Amsterdam. This paper also benefited from feedback we received on its *descendants* at the Central APA 2012, the University of Maryland, Georgetown University and Kai von Fintel and Sabine Iatridou's Fall 2012 seminar at MIT. Part of this work was carried out while the second author was a visiting researcher at Northwestern University, whose hospitality is gratefully acknowledged. The submission history of this paper is as follows: First version submitted March 28, 2011; first decision rendered April 29, 2012; revised version submitted October 16, 2012; accepted March 04, 2013.

References

Barker, C. 2010. Free choice permission as resource sensitive reasoning. *Semantics & Pragmatics*, 3(10):1–38.

- Belnap, N., M. Perloff, and M. Xu. 2001. *Facing the Future*. Oxford University Press.
- Cantwell, J. 2008. Changing the modal context. *Theoria*, 74:331–351.
- Cariani, F. 2009. *The Semantics of ‘ought’ and the Unity of Modal Discourse*. PhD thesis, UC Berkeley.
- Cariani, F. forthcoming. *Ought* and resolution semantics. forthcoming in *Noûs*.
- Carr, J. ms. Subjective ought. MIT manuscript.
- Charlow, N. forthcoming. What we know and what to do. forthcoming in *Synthese*.
- DeRose, K. 1991. Epistemic possibilities. *The Philosophical Review*, 100:581–605.
- Egan, A. 2007. Epistemic modals, relativism, and assertion. *Philosophical Studies*, 133:1–22.
- Egan, A., J. Hawthorne, and B. Weatherson. 2005. Epistemic modals in context. In Preyer, G. and P. Peter, editors, *Contextualism in Philosophy*. Oxford University Press.
- Egan, A. and B. Weatherson, editors. 2011. *Epistemic Modality*. Oxford University Press, Oxford.
- Fagin, R., J. Y. Halpern, Y. Moses, and M. Y. Vardi. 1995. *Reasoning about Knowledge*. MIT Press.
- von Fintel, K. and A. S. Gillies. 2010. Must . . . stay . . . strong! *Journal of Semantics*, 18:351–383.
- von Fintel, K. and A. Gillies. 2007. An opinionated guide to epistemic modality. *Oxford Studies in Epistemology*, 2:32–62.
- von Fintel, K. and A. Gillies. 2008. CIA leaks. *Philosophical Review*, 117:77–98.
- von Fintel, K. ms. The best we can (expect to) get? challenges to the classic semantics for deontic modals. presented at the 2012 Central APA, Chicago, IL. <http://mit.edu/fintel/fintel-2012-apa-ought.pdf>.
- von Fintel, K. and S. Iatridou. ms. What to do if you want to go to harlem? MIT.
- Frank, A. 1996. *Context Dependence in Modal Constructions*. PhD thesis, Universität Stuttgart.
- Geurts, B. 1999. *Presuppositions and Pronouns*. Elsevier, Oxford.
- Geurts, B. 2005. Entertaining alternatives. *Natural Language Semantics*, 13:383–410.
- Geurts, B. ms. On an ambiguity in quantified conditionals. manuscript, University Njemegen.
- Gillies, A. 2010. Iffiness. *Semantics and Pragmatics*, 3:1–42.
- Hacking, I. 1967. Possibility. *The Philosophical Review*, 76:143–168.
- Hamblin, C. L. 1958. Questions. *Australasian Journal of Philosophy*, 36:159–168.
- Horty, J. 2001. *Agency and Deontic Logic*. Oxford University Press.
- Horty, J. and N. Belnap. 1995. The deliberative *Stit*: a study in the logic of action, omission, ability and obligation. *Journal of Philosophical Logic*, 24:583–644.
- Jackson, F. 1991. Decision theoretic consequentialism and the nearest dearest objection. *Ethics*, 101(3):461–482.
- Kaufmann, S., C. Condoravdi, and V. Harizanov. 2006. Formal approaches to modality. In Frawley, W., editor, *The Expression of Modality*, pages 71–106. Mouton de Gruyter.

- Kaufmann, S. and M. Schwager. 2011. A unified analysis of conditional imperatives. In *Proceedings of SALT 19*, pages 223–238. CLC Publications.
- Kolodny, N. and J. MacFarlane. 2010. Ifs and Oughts. *Journal of Philosophy*, 107: 115–143.
- Kratzer, A. 1977. What ‘must’ and ‘can’ must and can mean. *Linguistics and Philosophy*, 3(1):337–355.
- Kratzer, A. 1978. *Semantik der Rede. Kontexttheorie, Modalwörter, Konditionalsätze*. Scriptor, Königstein.
- Kratzer, A. 1981. The notional category of modality. In Eikmeyer, H.-J. and H. Rieser, editors, *Words, Worlds, and Contexts: New Approaches in World Semantics*. de Gruyter, Berlin.
- Kratzer, A. 1986. Conditionals. In Farley, A. M., P. T. Farley, and K.-E. McCullough, editors, *Proceedings of CLS 22, Part 2: Papers from the Parasession on Pragmatics and Grammatical Theory*, pages 1–15.
- Kratzer, A. 1991. Modality. In von Stechow, A. and D. Wunderlich, editors, *Semantik: Ein internationales Handbuch der zeitgenössischen Forschung. [=Semantics]*, pages 639–650. de Gruyter.
- Kratzer, A. 2012. *Modals and Conditionals*. Oxford University Press.
- Lewis, D. 1973. *Counterfactuals*. Harvard University Press.
- Lycan, W. 1993. MPP, RIP. *Philosophical Perspectives*, 7:411–428.
- Lycan, W. 2001. *Real Conditionals*. Oxford University Press.
- MacFarlane, J. 2011. Epistemic modals are assessment-sensitive. In Egan and Weatherson (2011).
- Mellor, D. H., editor. 1990. *Philosophical Papers: F. P. Ramsey*. Cambridge University Press.
- Parfit, D. 1988. What we together do. manuscript, Oxford University.
- Parfit, D. 2011. *On What Matters*. Oxford University Press.
- Portner, P. 2009. *Modality*. Oxford University Press, Oxford.
- Ramsey, F. P. 1929. General propositions and causality. Printed in Mellor (1990), pages 145–163.
- Sæbø, K. J. 2002. Necessary conditions in a natural language. In Fery, C. and W. Sternefeld, editors, *Audiatur Vox Sapientiae. A Festschrift for Arnim von Stechow*. Akademie Verlag, Berlin.
- Schroeder, M. 2010. Oughts, agents and actions. *The Philosophical Review*, 120(1): 1–41.
- Schwager, M. 2006. Conditionalized imperatives. In Tancredi, C., M. Kanazawa, I. Imani, and K. Kusumoto, editors, *Proceedings of SALT XVI*. CLC Publications, Ithaca, NY.
- Simons, M. 2007. Semantics and pragmatics in the interpretation of *or*. *Proceedings of SALT 15*.
- Stephenson, T. 2007. Judge dependence, epistemic modals, and predicates of personal taste. *Linguistics and Philosophy*, 30:487–525.
- Thomason, R. H. 1981. Deontic logic as founded on tense logic. In Hilpinen, R., editor, *New Studies in Deontic Logic*, pages 165–176. Reidel.
- Yalcin, S. 2007. Epistemic modals. *Mind*, 116(4):983–1027.

- Zimmermann, T. E. 2000. Free choice disjunctions and epistemic possibility. *Natural Language Semantics*, 8:255–290.
- Zvolenszky, Z. 2002. Is a possible-worlds semantics of modality possible? A problem for Kratzer's semantics. In Jackson, B., editor, *Proceedings of SALT XII*.