Deliberative Modality under Epistemic Uncertainty

Fabrizio Cariani, Stefan Kaufmann & Magdalena Schwager

March 2011

---

1 Introduction

The modalized sentence in (1) is easily understood as expressing that reading said article is the optimal 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 individuated and ranked according to certain contextual parameters (Kratzer, 1977, 1981). 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 stake in practical deliberations like (1) ***deliberative modality***, and we will assume that, together with at least deontic (in the narrow sense, that is, related to rules and laws), bouletic (related to wishes) it belongs to the so-called prioritizing modal flavors (in contrast to epistemic and dynamic ones, cf. Portner 2009).[1]

The ability of this approach to provide an adequate representation 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):[2]

> 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

[1] Deliberative modality may be closely related to teleological, that is, goal oriented, modality. Yet, teleological modality is sometimes claimed to involve an additional 'in order to'-parameter (a ***designated goal***, cf. von Fintel and Iatridou 2005). As we do not want to go into this discussion, and no such parameter will be needed for the cases of practical deliberation which 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).

[2] The 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).

water, and just one miner, the lowest in the shaft, will be killed.

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 judgements:

(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 judgements strikes us as correct—provided that we stick to the deliberative interpretation of *ought*. In principle, it would be best to block the shaft in which the miners are, thus (2b) and (2c) seem to be true. But given that we don't know and have no way of finding out,[3] it is best to block neither shaft and thus guarantee that we save nine of the ten. So, (2a) seems to be true as well. However, from (3), which is undisputedly true in the scenario under consideration, together with (2b) and (2c) 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 (2a). And, indeed, (4) doesn't seem to be a correct outcome of practical reasoning.

The puzzle poses two distinct problems:

LOGICAL How should we block the derivation of (4) from (2b), (2c) and (3)?
SEMANTIC How can our semantic machinery derive all of the salient verdicts?

Kolodny and MacFarlane (2010) discuss various strategies to solve the logical problem—by rejecting one or more of the premises or giving up on one or more of

[3] It is vital to Kolodny and MacFarlane's (2010) reasoning that the decision has to be taken without there being a chance of finding out where the miners are.

the inference rules. They end up jettisoning modus ponens; without this rule, they argue, from (2b), (2c), (3) we can no longer derive (4). As they acknowledge, this leaves open the question concerning the correct semantics.

In much recent research in natural language semantics, neither the validity of modus ponens for indicative conditionals nor a classical analysis of disjunction is taken for granted. In light of this, we might expect that, under standard assumptions in linguistic semantics, the semantic problem would not arise. This expectation turns out to be false. In Section 2, we show that the problem arises given any straightforward specification of the general assumptions in the standard framework of possible worlds semantics. In Section 3, we propose an extension of the standard framework and show how this accounts for 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. Section 3). At the same time, we aim to improve on their proposed answer to the semantic problem in two respects: First, we will claim that our approach gives a better diagnosis of the puzzle. Second, our approach produces more systematic explanations of how the truth-values of deliberative modalities evolve as context evolves. Finally, our proposal constitutes a conservative extension of the standard analysis (which can be simulated as a special case of our approach, as we show in Section 4).

2 The miners problem in Kratzer's framework

Kolodny and MacFarlane's solution to the miners problem relies on the rejection of modus ponens. The standard semantics for conditionals in the Kratzer-style framework for modals and conditionals accounts for failures of modus ponens as observed in non-practical inference as well. In this section, we briefly introduce the framework and evaluate its predictions for the miners problem. We argue that the framework 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

In this section we briefly introduce the formal tools standardly adopted in linguistic semantics. We take as given a non-empty set W of possible worlds. Propositions are represented as subsets of W. Atomic sentences denote propositions. We first define the interpretation of modal verbs and then discuss various possibilities of applying the framework to the sentences in the miners example.

Following Kratzer (1981), modal sentences are interpreted relative to two parameters f and g. Both are ***conversational backgrounds*** in her terms, i.e., functions from worlds to sets of propositions. The first 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. It is called the modal base. 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.[4] The second parameter g is called an ***ordering source***. For each world w, it induces a pre-order on the set of worlds. Depending on the way g is interpreted, $u \leq_{g(w)} v$ may mean that u is less far-fetched than v, preferable to v, or "better" than v in some other contextually salient respect. Formally, the relation is defined as follows, for each $u, v \in W$:

(5) $u \leq_{g(w)} v$ iff $\{p \in g(w) | v \in p\} \subseteq \{p \in g(w) | u \in p\}$

We add a ***human necessity*** modal operator $\boxplus$ to the language whose interpretation relative to a world w, modal base f and ordering source g is defined in (6a). A special case is the (familiar) ***simple necessity*** operator $\Box$, defined in (6b) as shorthand for human necessity relative to a constant ordering source whose value is the empty set of propositions. In this case, the sentence is true iff its prejacent is true at all worlds in the modal base.

(6) a. $\boxplus\varphi$ is true at $\langle w, f, g\rangle$ iff
$\forall v \in \bigcap f(w)\Big[\exists u \in \bigcap f(w)\Big[u \leq_{g(w)} v \wedge \forall z \in \bigcap f(w)\Big[z \leq_{g(w)} u \rightarrow z \in \varphi\Big]\Big]\Big]$.

b. $\Box\varphi$ is true at $\langle w, f\rangle$ iff $\boxplus\varphi$ is true at $\langle w, f, \lambda v.\emptyset\rangle$.

If the ordering source $g(w)$ is finite, there is guaranteed to be a set of minimal or "best" worlds. We assume for simplicity that this is the case (Lewis 1973's ***Limit Assumption***) and refer to this set as in (7a). Under the given assumptions, the semantics for $\boxplus$ can equivalently be spelt out as in (7b).

(7) a. $O(w, f, g) =_{\text{df}} \left\{u \in \bigcap f(w) \,\middle|\, \forall v \in \bigcap f(w)\Big[v \leq_{g(w)} u \rightarrow u \leq_{g(w)} v\Big]\right\}$

b. $\boxplus\varphi$ is true at $\langle w, f, g\rangle$ iff φ is true at all worlds in $O(w, f, g)$.

Nothing in the given scenario seems to turn on the question whether the ordering source is finite or not. We will therefore use the simpler definition in (6a) without loss of generality.

[4] There 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 $\bigcap f(w)$.

As mentioned above, in Kratzer's system the interpretation of a modal crucially depends on two contextual parameters, the modal base and the ordering source. According to the received view, *ought* expresses human necessity relative to these parameters. What are their most plausible settings in the miners case?

2.2 Ordering source

The standard assumption about the ordering source for *ought* is that it represents our goals or moral obligations. In principle, these two notions are distinct and give rise to ***teleological*** and ***deontic*** modal flavors. In the given scenario, our intuitions may waver as to how the particular ordering source should be labeled. But assuming that the moral obligations we are under are also goals we strive to attain, we need not worry too much about which of the two notions is at play. What is important in the given situation 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."[5]

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 (8).

(8) $g(w)$ = {all miners are saved, at least 9 miners are saved, ...,
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 semantics, which we will later compare to Kolodny and MacFarlane's proposal.

Before we start, notice that the set of possible words compatible with the description of the scenario can be split into six mutually exclusive subsets, each of

[5]Note 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 (2a)-(2c).

Table 1: Classes of worlds in the given scenario

cell	*position of miners*	*action*	*miners lost*
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

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

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.[6] 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.

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. That 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 it is.

[6] For some general reflection on this issue, cf. Section 3.3.

Table 2: Parameters of evaluation with circumstantial modal base

	w_A	w_B
modal background	$\bigcap f_c(w_A)$ = AA ∪ AB ∪ AN	$\bigcap f_c(w_B)$ = BA ∪ BB ∪ BN
ordering source	$g(w_A) = g(w) = g(w_B)$	
ranking[7]	AA $<_{g(w_A)}$ AN $<_{g(w_A)}$ AB	BB $<_{g(w_B)}$ BN $<_{g(w_B)}$ BA
minimal worlds	$O(f_c, g, w_A)$ = AA	$O(f_c, g, w_B)$ = BB

Put differently, the possible worlds that are consistent with the description of the scenario fall into two classes according as the miners are trapped in shaft A or B. In Table 1, they are the sets AA ∪ AB ∪ AN and BA ∪ BB ∪ BN, respectively. Let w_A and w_B be two arbitrary worlds from each of these two classes. The world w of evaluation – the world in which the described agents actually carry out their deliberation as to which *ought*s are true – is similar to exactly one of them in all relevant respects. Thus the value of the circumstantial modal base f_c at w is one of the two sets of propositions in (9). The description does not specifiy which one it is.

(9) $f_c(w) \in \{f_c(w_A), f_c(w_B)\}$

a. $f_c(w_A)$ = {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. $f_c(w_B)$ = {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 the sentence (2a) at either of these worlds are as given in Table 2. Recall that the sentence is true with these settings at a world w if and only if all worlds in the set of minimal worlds $O(f_c, g, w)$

[7]Here and below, expressions of the form '$X\rho Y$', where X, Y are subsets of the domain and range of ρ, respectively, are shorthand for the statement that $x\rho y$ for all $x \in X, y \in Y$.

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 (2a) is false on the circumstantial reading. In contrast, (4) is true at both w_A and w_B.[8]

(2a) We ought to block neither shaft. ✗

(4) Either we ought to block A or we ought to block B. ✓

These predictions are squarely at odds with Kolodny and MacFarlane's judgments about the miners scenario. While they do not deny the existence of an objective reading for *ought*, they argue that this reading does not play a role in deliberation, simply because the agents are not in an epistemic position to evaluate *ought*-sentences in this way.

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 (10).

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

Here, too, our intuition is that while it may be judged false in view of the deliberative process (in particular, it is at odds with (2a)), it can also be understood as true in the objective interpretation. The combination of a circumstantial modal base and a deontic ordering source predicts (10) to be true.

Putting that issue aside, we grant that the reading Kolodny and MacFarlane have in mind (on which (2a) 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 do justice to this deliberative reading, we need to keep looking for the right combination of parameters to account for it.[9]

[8] Here and in the following, ✗ and ✓ indicate the predictions of the respective theory under discussion, rather than speakers' intuitions.

[9] 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 rely on making all sentences (2a)–(4) compatible by postulating

The problem with the circumstantial modal base is that it takes into account the actual location of the miners. 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 (in the given scenario, the group formed by the speaker and an unspecified set of further individuals).

2.3.2 Epistemic

This would suggest an ***epistemic*** modal base f_e whose value for the actual world w is specified by Kolodny and MacFarlane's description as the set of propositions in (11). Notice that this epistemic modal base, unlike the circumstantial one above, is the same at all worlds consistent with the description.

(11) $f_e(w) = f_e(w_A) = f_e(w_B)$
= {there are two shafts A and B,
10 miners are either in shaft A or in shaft B,
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}

Consider how the sentences in question are evaluated relative to this modal base. The relevant parameter settings are shown in Table 3. The set of minimal worlds under the induced ordering is AA ∪ BB; at all these worlds, we block a shaft. So the sentence (2a) is again not true. In this regard, the move from a circumstantial reading to an epistemic one did not help matters.

(2a) We ought to block neither shaft. ✗

(4) Either we ought to block A or we ought to block B. ✗

different readings for the occurrences of *ought*. It is important to note, however, that this objection is orthogonal to our concerns. Although we believe that (2a) 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 this. Rather, our point is that the truth of (2a) cannot be captured by the standardly assumed settings for modal base and ordering source and, as we will argue below, requires a modification of the Kratzer-style framework.

Note that Kolodny and MacFarlane (2010) raise another issue for context dependence: They argue that variability along the lines discussed here (for their account, whose information state is relevant) should not be captured as dependence on the context of utterance, but on the one of assessment. As our solution could easily be translated under either assumption, we ignore this issue and stick with the more standard notion of context dependence.

Table 3: Classes of worlds in the epistemic modal base

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

However, the epistemic modal base does yield a different prediction about (4): Since relative to the epistemic modal base 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. Insofar as epistemic modal bases are naturally associated with deliberative (as opposed to objective) *ought* this prediction, this result is desirable. [10] Similarly, as the minimal worlds are divided as to which of the two shafts is blocked, (10) is now predicted to be false: There is no shaft such that we block it at all minimal worlds.

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

To sum up, the predictions resulting from an epistemic modal base do not fit our intuitions with respect to deliberative *ought*, under which (2a) should be true. Nor does it capture the objectivist judgment that (2a) is false whereas (4) and (10) are true. Recall that these were captured correctly by the circumstantial modal base. Thus the epistemic modal base fails to deliver a Kratzer-style solution to the miners problem.

At this point one might think that we just happened to pick the wrong epistemic modal base by relying on the one that specified 'our' (the deliberating agents in Kolodny and MacFarlane's scenario) information. In principle, an epistemic modal base might correspond to any other agent's (and possibly an even more abstract) body of information.[11] But even if some such other epistemic modal base were

[10] Some theories of the semantics and pragmatics of English *or* may be able to derive (4) as true with respect to an epistemic modal base (consider in particular Geurts, 2005). However, the complexities of such a treatment are too far removed from the main concerns of this paper to warrant a detailed discussion here. Also, examples like (10), or its narrow-scope analogue:

(i) We ought to block at least one shaft.

prove that problems can arise even in cases that do not involve *or*.

[11] This flexibility has recently been widely discussed in the literature on epistemic modals (DeRose, 1991; Egan, 2007; von Fintel and Gillies, 2007, 2008; Hacking, 1967; MacFarlane, to

salient, this would not help: Any body of information that specifies the location of the miners behaves like the circumstantial modal base in predicting (2a) to be false because at all best worlds that particular shaft gets blocked; any body of information that does not specify the location of the miners (like 'our' epistemic modal base) predicts it to be false because the best worlds are divided into worlds where A is blocked and worlds where B is blocked.

2.3.3 Deliberative *ought* in conditionals

The conditionals in (2b) and (2c) play an important role in Kolodny and MacFarlane's argument.

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

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

Kratzer's framework gives considerable leeway in the way it is applied to conditionals and easily predicts both conditionals to be true in the miners scenario, under all plausible settings of the parameters for this situation.

In Kratzer's formalization, *if*-clauses add the proposition denoted by the antecedent to the modal base. We adopt Kolodny and MacFarlane's notation and refer to the result of modifying a modal operator $\Box$ with a conditional antecedent *if* φ as '$[\textsc{if}\,\varphi]\,\Box\psi$'. To simplify the exposition, we introduce the following notational shorthand: For all modal bases f and propositions p, the *update of* f *with* p is a function $[f{+}p]$ from worlds to sets of propositions such that $[f{+}p](w) = f(w)\cup\{p\}$ for all worlds w. It is easy to check that this results in a restriction of the modal background $\bigcap f(w)$ to p-worlds.[12,13]

(12) $[\textsc{if}\,\varphi]\,\Box\psi$ is true at $\langle w, f, g\rangle$ iff $\Box\psi$ is true at $\langle w, [f{+}[\![\varphi]\!]], g\rangle$
iff for all $w' \in O(w, [f{+}[\![\varphi]\!]], g)$, ψ is true at w'.

When no overt modal is present in the sentence, it is generally assumed that a covert epistemic modal is filled into the logical form, and the antecedent serves to restrict the modal base of that epistemic modal. Opinions diverge, however, on what happens when the only overt modal in a conditional is deontic, as in (13).[14] In

appear; Stephenson, 2007, among others).

[12] $\bigcap[f{+}p](w) = \bigcap(f(w)\cup\{p\}) = (\bigcap f(w))\cap p$.

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

[14] Here 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.

principle, there are two possible construals for this case: Either the deontic modal is itself restricted by the conditional antecedent, as in (13a), or the antecedent restricts a covert epistemic modal and the deontic modal is part of the consequent (13b). Schwager (2006) labels these two interpretations the ***Overt*** and ***Covert Conditional Operator*** contrual, respectively.

(13) If φ, ought ψ.

a. $[\textsc{if}\ \varphi]\ \boxplus_d \psi$ — Overt Conditional Operator (OCO)

b. $[\textsc{if}\ \varphi]\ \boxplus_e [\boxplus_d \psi]$ — Covert Conditional Operator (CCO)

Both construals have some popularity, and both face certain problems. The current tendency seems to be that the construal in (13b) is needed at least for some examples, and that it may well be the only construal needed. A detailed discussion of the relevant arguments would lead us too far afield (see Frank, 1996; Zvolenszky, 2002; Kaufmann and Schwager, t.a., among others). Nor do we need to not take a stance on these matters in this paper. We just want to point out that satisfactory predictions for the conditonals in the miners scenario do not depend on a particular setting of these parameters.

Instantiated with sentence (2b), the two construals correspond to (14a) and (14b), respectively (here '$\Box_e$' stands for the silent epistemic operator).

(14) If the miners are in shaft A, we ought to close shaft A.

a. $[\textsc{if}\ [\![\text{the miners are in A}]\!]]\ \text{ought}[\![\text{we close A}]\!]$ (OCO)

b. $[\textsc{if}\ [\![\text{the miners are in A}]\!]]\ \Box_e[\text{ought}[\![\text{we close A}]\!]]$ (CCO)

It is now easy to see that under both construals and under all plausible assumptions about the choices for the modal bases, the conditionals are predicted to be true in the miners scenario. This is because under all plausible settings, the following holds: (i) the modal background relative to which *ought* ends up being interpreted consists entirely of antecedent-worlds; (ii) the prejacent of *ought* is consistent with this modal background; and (iii) nothing in the description of the scenario motivates a change to a different ordering source.

Consider first the OCO 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, the consequent *ought* ψ is interpreted relative to $\bigcap [f+[\![\text{the miners are in A}]\!]]\,(w)$. As the ordering source is unchanged, these worlds are still ranked according to how many of the miners are saved, which singles out as best those worlds where 'we' block shaft A. Similarly for (2c) and

shaft B.[15]

Under the CCO construal, the consequent is interpreted pointwise at worlds at which the antecedent is true (as before, let w_A be an arbitrary such world), and relative to a modal base of its own that is not necessarily that of the covert epistemic operator. As before in Section 2.3, there are multiple possibilities, and again, all solutions are perfectly parallel for (2b) and (2c).

The first option is a circumstantial modal base $f_c(w_A)$. As before, the location of the miners is among the 'relevant facts' on any possible interpretation of this conversational background, hence all these worlds in this background are ones at which they are in A. Assuming further that 'we' know what our preferences are (i.e. for all $w', w'' \in \bigcap f_e(w) : g(w') = g(w'')$) and that those preferences call for saving as many miners as possible, the best ranked worlds are all worlds at which we block shaft A.

The second option would be an epistemic modal base for *ought*. Notice that under this construal we are really dealing with two values of the modal base: One, $f_e(w)$, for the interpretation of the covert conditional operator, and one, $f_e(w_A)$, for the interpretation of *ought* at worlds w_A in $\bigcap [f_e{+}[\![A]\!]](w)$. In principle, the two are independent and nothing ensures that $f_e(w_A)$ "inherits" the restriction to antecedent-worlds from $[f_e{+}[\![A]\!]](w)$.[16] Intuitively, however, *ought* can be interpreted relative to the restricted modal base.[17] The standard solution for such cases under the CCO construal is to assume a mechanism akin to anaphora resolution that makes the restriction of the outer modal to the antecedent-worlds carry over to the inner modal as well. That is, the modal base of *ought* would also be restricted by *the miners are in shaft A*, and the reasoning then proceeds as for the unembedded case. Frank (1996) and Geurts (1999) offer independent motivation for the anaphoric nature of modal bases.[18]

[15]There are some "pathological" cases among the possible OCO-readings. For instance, as we argued above, a circumstantial modal base $f_c(w)$ should be assumed to contain the actual location of the miners at w. But then, the addition of the antecedent of (14a) 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.

[16]Indeed, assuming that f_e is introspective, we would expect $\bigcap f_e(w_A) = \bigcap f_e(w)$.

[17]We do not claim that *ought* only has the restricted reading, though. For instance, the following is true, albeit implausible in the scenario:

(i) 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.

[18]One may wonder about the possibility of epistemic modal bases representing belief states or information sources other than that of the covert epistemic operator. We do not deny the existence of such readings for *ought* in conditional consequents, but we refrain from exploring them in detail

To sum up, while one may have independent reasons for preferring some of the mentioned possibilities over others, they all interpret *ought* in the consequent of (2b) relative to a set of worlds in which the miners are in shaft A (or in shaft B, for (2c)), hence all render the conditionals true relative to the deontic ordering source we introduced above.

2.4 Introducing information dependence

Let's take stock: We have convinced ourselves that the Kratzer-style framework predicts the conditionals (2b) and (2c) to be true under all plausible parameter settings. In contrast, on any setting for the modal base and ordering source standardly considered, the framework fails to predict the deliberative reading on which (2a) is true in the given scenario.

In many respects, Kolodny and MacFarlane's (2010) treatment of modal verbs is similar to Kratzer's, but it has no problems in accounting for the truth of deliberative *ought*. A brief comparison of the two accounts will help us detect the crucial limitation in Kratzer's account: In contrast to Kolodny and MacFarlane's (2010) treatment of modals, it fails to treat them as ***seriously information dependent***. 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 will ultimately develop a variant of the Kratzer-style framework that embraces serious information dependence.

2.4.1 Comparison with information dependent *ought*

Like Kratzer, Kolodny and MacFarlane take the interpretation of *ought* at a given world w to be sensitive to two parameters.[19] The first is an ***information state***, formally represented as a set i of worlds. But for at least some modals (notably deontic ones), not all worlds in i are quantified over. The second parameter is a ***selection function*** mapping i to a set of worlds which then constitutes the domain of quantification for the interpretation of the modal operator.

Both epistemic and deontic modals are sensitive to the information state, but they differ in the selection function. An ***epistemic*** selection function e trivially selects the input 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

here. The scenario in question does not seem to rely on them, and a detailed investigation would take us too far afield.

[19] They 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 the linguistic literature.

dub "realism" and spell out as the condition that $d(i) \subseteq i$ for all i.[20]

Thus far there appears to be no semantically significant difference between Kolodny and MacFarlane's and Kratzer's framework. Conceptually the information state in the former corresponds to the modal base in the latter, while Kolodny and MacFarlane's selection function corresponds to the ordering source $g(w)$: It singles out the domain of worlds over which the truth definitions for the modal operators quantify:[21] The basic skeleton of the truth conditions is quite similar. Restricting our attention for the sake of simplicity to modal sentences $\ulcorner \Box\varphi \urcorner$ whose prejacent φ does not itself contain any modals, we can state the truth conditions as follows (the subscript 's' on the modal operator in (15a) stands for the selection function):[22]

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

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 information state (since $d(i) \subseteq i$ and $e(i) = i$) or the modal background (since $O(w, f, g) \subseteq \bigcap f(w)$ for any g). Clearly, then, the crucial difference that explains why Kratzer's, but not Kolodny and MacFarlane's, framework runs into problems with deliberative *ought* must lie in the way in which the relevant subset is identified.

2.4.2 Serious information dependence

Both frameworks under discussion interpret *ought* relative to a set of deontically "best" worlds selected from a modal background that is characterized by some body of information. But crucially, Kratzer's is committed to a particular mechanism for picking out the best worlds: This happens via an ordering source the value of which has to correspond to a contextually salient set of preferences or

[20] Note that this is different from Kratzer's (1981) notion of realism, which in Kolodny and MacFarlane's framework would amount to the requirement that i contain the world of evaluation. Instead, the intention here is that only worlds are selected that might, as far as the information in i is concerned, be the actual one.

[21] Kolodny and MacFarlane implicitly make the Limit Assumption and do not provide for the possibility that it might be violated. In this sense, Kratzer's framework is more general. However, this difference seems to be irrelevant to the ability of the two frameworks to explain the data at hand.

[22] If the prejacent contains a modal operator, its truth is not determined by the world w' alone. Kratzer originally assumed, similarly to Kolodny and MacFarlane, that the modal base relevant for φ in this case is f. But this assumption is no longer widely accepted in the linguistic literature. Instead, the modal base for φ may be identified with f via co-reference, but may also be set to some other contextually given value.

rules. No such requirement is imposed by Kolodny and MacFarlane. They take the selection function to be primitive and make no attempt to predict its value for a particular information state i from any property of that information state, as long as $d(i) \subseteq i$ (the requirement that d be "realistic" as discussed above). Moreover, as evidenced by the miners scenario, with deliberative modals, what counts as the 'best' course of action could change as more information becomes available. Crucially, this can happen even if the new information does not eliminate any of the options for action. Upon learning that the miners are in shaft A, blocking shaft A would constitute the best course of action, even though blocking neither shaft (the optimal course of action so far) remains an epistemic possibility. In Kolodny and MacFarlane's framework it is obvious that "more information" corresponds to evaluation with respect to a subset of the information state considered previously.

To allow for the relative "goodness" of various choices of action to flip in moving to a richer information state, they explicitly require that the deontic selection function upon which the interpretation of deliberative *ought* relies have a property which they call ***serious information dependence***.

Definition 1 (Serious information dependence – Kolodny and MacFarlane)
A deontic selection function d is seriously information-dependent *iff for some $i_1, i_2 \subseteq i_1$, there is a world $w \in i_2$ such that $w \in d(i_1)$ but $w \notin d(i_2)$.*

A direct comparison with Kratzer's account requires some care. So far we have informally put various pieces of the two frameworks side by side; however, while this loose talk is intuitively perfectly sensible, there is strictly speaking no one-to-one correspondence between the various ingredients: Kolodny and MacFarlane's i and f are of different types, and so are the selection function d and the ordering source g. We can, however, easily define an auxiliary notion of "Kratzer selection function" in terms of g 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 depends only on the modal background, a set of worlds:

Fact 1
For all possible worlds w, modal bases f, ordering sources g and propositions p, if $\bigcap f(w) = p$, then $O(w, f, g) = O(w, \lambda v.\{p\}, g)$.

This follows immediately from (7a) above, where $O(w, f, g)$ was defined in terms of the order $\leq_{g(w)}$ without any reference to worlds not in $\bigcap f(w)$. We can now define, for any g and w, a function mapping sets of worlds 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 sets of worlds to sets of worlds such that for all $p \subseteq W, d_{g,w}(p) =$*

$O(w, \lambda v.\{p\}, g)$.

Theorem 1

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 some $i_1, i_2 \subseteq i_1$ and world $v \in i_2$ such that (i) $v \in d_{g,w}(i_1)$ and (ii) $v \notin d_{g,w}(i_2)$. By (ii), there is a world $u \in i_2$ such that $u <_{g(w)} v$. But $u \in i_1 \supseteq i_2$, hence $v \notin d_{g,w}(i_1)$, contradicting (i). □

Theorem 1 brings out the crux of the problem with the Kratzer-style semantics. With a seriously information-dependent selection function, it is possible for the relative "goodness" of worlds to flip in moving to a richer information state: Some worlds may start out as ideal relative to the information in $f(w)$ and still be consistent with, but no longer ideal relative to that in $[f+[\![A]\!]](w)$. Kratzer semantics lacks the resources to handle this feature of deliberative modality.

This can be seen clearly in the miners scenario: Even if we could find values for modal base f and ordering source g such that, at the world of evaluation w, (2a) is true with respect to these parameters, it is impossible to render the sentence false at a world w' which differs from w only in that $f(w')$ contains more information than $f(w)$ and hence $\bigcap f(w') \subset \bigcap f(w)$, while $g(w) = g(w')$. Clearly, this is at odds with our intuitions about the miners scenario.

The result has further interesting implications for the interpretation of the conditionals. As long as the conditional antecedent does not contain modals (as in (2b)-(2c)), Kolodny and MacFarlane settle for an interpretation on which conditional antecedents restrict the modal base, the consequent is evaluated with respect to that restricted information state.[23] Given this stance, it is crucial that *ought* be seriously information dependent in order for (2a), (2b) and (2c) to be true with respect to one and the same point of evaluation.

The limitation of Kratzer's framework carries over to its application to conditionals: If conditionals are interpreted in such a way that the modal background of *ought* in the consequent is restricted to a subset of the modal background in the unembedded case, then (2a) and either of the conditionals (2b) and (2c) cannot be true at the same world with *ought* being evaluated relative to the same g. If we reconsider briefly the construals that tend to be adopted within the Kratzer-style framework, it becomes immediately obvious that, in the given scenario, the OCO-construal does not allow for the simultaneous truth of (2a) and either of the two conditionals, regardless of the modal base (and ordering source) chosen.

[23] In the general case, they adopt a more complicated account inspired by Yalcin (2007).

Things are somewhat more intricate with the CCO-construal: Remember that here, the modal base of the deontic modal in the consequent is evaluated pointwise at the worlds in the modal background of the outer modal which we have assumed to express simple epistemic necessity. If we assume that *ought* is interpreted with respect to the same epistemic modal base and inherits the restriction to the antecedent worlds, its occurrence in the conditional is clearly evaluated at a subset of the modal background relevant for (2a). If, instead, *ought* is interpreted with respect to a circumstantial modal base both in (2a) and in the conditional antecedent, there is little reason based on the scenario to assume that there should be a difference between the relevant circumstantial modal backgrounds at the "actual" world at which (2a) is interpreted on the one hand, and at those epistemically accessible antecedent worlds that agree with the actual world with respect to the location of the miners. Hence, as one of the conditionals agrees with the unembedded *ought* on the modal background for the occurrence of *ought* in its consquent, it is impossible that that conditional could single out as optimal a different action from the one singled out by the unembedded *ought* – as long as the ordering source is constant. Finally, *ought* may be interpreted with an epistemic modal base in (2a) but with a circumstantial one at each of the epistemically accessible antecedent-worlds in the conditional. In this case, since we may assume that in the miners scenario the circumstantial modal background is a subset of the epistemic one and the same for all antecedent-worlds, we once again conclude that *ought* in the conditional is interpreted relative to a subset of the modal background relative to which the unembedded *ought* is interpreted, hence the failure of serious information dependence becomes again relevant.

Things look bleak for Kratzer's framework at this point. We saw above that it provides no obvious way to account for the truth of (2a) in the scenario. We further found that due to its lack of serious information dependence, any combination of the parameters of interpretation which (i) interprets *ought* in the conditional consequent relative to a modal background that is a subset of the one for (2a) and (ii) renders the conditionals true, will have trouble accounting for the truth of (2a) relative to the same ordering source. But all the most sensible parameter settings do have both properties (i) and (ii). One way to proceed might be to seek justifications for more far-fetched parameter settings which avoid the subset relation between the two modal backgrounds. But if it is hard to come up with such parameter settings, it is even harder to think of a credible justification for them in the scenario.[24] We

[24]One may wonder about the possibility of epistemic modal bases representing belief states or information sources other than that of the covert epistemic operator. We do not deny the existence of such readings for *ought* in conditional consequents, but we refrain from exploring them in detail here. The scenario in question does not seem to rely on them, and a detailed investigation falls outside of the scope of our argument here.

Table 4: Selection function in the miners scenario

Information state	*Best Worlds*
i = AA ∪ AB ∪ AN ∪ BA ∪ BB ∪ BN	$d(i) \subseteq$ AN ∪ BN
i_A = AA ∪ AB ∪ AN	$d(i_A) \subseteq$ AA
i_B = BA ∪ BB ∪ BN	$d(i_B) \subseteq$ BB

think this is hardly a promising direction. But this should not be taken as a reason to abandon the core virtue of Kratzer's framework. We take this virtue to be that it can transparently derive the comparative evaluation of worlds based on properties of the world of evaluation (namely the prevailing preferences, goals and norms). It is time to illustrate this point in greater detail.

2.4.3 In favor of transparency

Kolodny and MacFarlane open the door to 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 in terms of "goodness" can vary as a function of 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 4. 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 serious 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 are lost. Consider this pair:

(16) If you are healthy, you ought to read that article.

(17) 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, as resulting from a single initial pair of modal base and ordering source. 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 (16) 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 (17), 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.[25]

The account we get from Kolodny and MacFarlane is nowhere 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 (16) and (17).

3 Extending the standard analysis

In Section 2.1 we showed that the standard Kratzer-style translation of *ought* as a weak necessity operator with an ordering source reflecting the agents' preferences fails to predict the truth of (2a), 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. Thus as long as the modal base provides worlds of both types, (2a) cannot be true. Indeed, intuitively blocking the shaft in which the miners are is the best possible course of events. But the agents cannot rely on this in choosing among the actions 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 there is no way to find out where the miners are, neither of these two actions is "safe," and the third choice of blocking neither shaft (where nine of the miners are saved) outranks both.

[25] Note 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.

While the agents' ignorance is a crucial factor in this reasoning, it is not sufficient to explain the truth of (2a) (switching to an epistemic modal base did not solve the problem, cf. Section 2.3.2). 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 (2a), 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.[26] 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. We further assume that each decision problem presents actions that are mutually exclusive relative to the relevant modal base.[27]

Decision problems thus understood need not partition the logical space, or even the set of epistemically possible worlds. For one thing, they need not be logically incompatible. There are worlds at which we block both shaft A and shaft B, but those worlds are not among the ones that are compatible with the description of the scenario. The actions are only incompatible given the situation; hence the above

[26]Note 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 boths shafts is an option, we ought to do that.

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

[27]Based 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 Belnap and Horty (1995); Belnap et al. (2001); Horty (2001).

requirement that the corresponding propositions be mutually exclusive relative to the modal base. Nor do the actions constituting a decision problem have to jointly exhaust all possibilities, even those that are epistemically accessible. There are some possible worlds at which 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 observation (either semantically or pragmatically), but adjudicating the issue is largely orthogonal to our aims here.[28] We will simply assume that while the contextually assigned decision problem need not constitute a partition with respect to the logical space, it has to determine a partition of the deliberative modal's modal base f.[29] We use '$[\delta|f]$ ' to refer to the function that results from combining the modal base f with the decision problem, which, at each world, returns a restriction of the decision problem to the worlds in the modal base:

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,*

$$[\delta|f](w) = \left\{\bigcap(f(w) \cup \{p\}) \middle| p \in \delta(w)\right\}$$

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. Notice that the empty set may be in $[\delta|f](w)$. This is not useful, but we do not take the trouble to rule it out because as we will see below, the presence of the empty set has no effect on the truth conditions we derive for deliberative modals.

Intuitively, the propositions in $[\delta|f](w)$ represent what is known about the decision problem, in the sense of both which choices are available and what the consequences of each choice are. Consider a very simple example: Grace needs to go to the post-office on a cloudy day. We can represent her options as in:

(18) {Grace goes to the post-office with her umbrella, Grace goes to the post-office without the umbrella, Grace stays at home}

In the 'unfiltered' sense *going to the post-office with her umbrella* consists of all and only those worlds in which Grace undertakes that course of action. In the

[28] Incidentally, the formal similarity between decision problems in our model and question denotations is striking. 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.

[29] This 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.

'filtered' sense, however, this action is represented only by those possible worlds that are compatible with the available information (e.g. words at which the current day is cloudy) and in which she goes to the post-office with her umbrella.

For purely deliberative *ought*, the decision problem can only be spelt out as fine-grainedly as the respective actions that are chooseable for the agent.[30] In particular, in the miners scenario, the natural decision problem is (19):

(19) δ = {we block shaft A, we block shaft B, we block neither shaft}

To take into account the fact that deliberative modality is influenced by the agents' limited knowledge, the modal base is constituted not by the set of relevant circumstances, but by what is known to the deliberating agent—that is, the modal base is f_e.[31] The ordering source $g(w)$ is again constituted by the criteria used in the deliberation, for all worlds w.

(20) $g(w)$ = {all 10 miners are saved, at least 9 miners are saved, ..., at least 1 miner is saved }

What is distinctive about our approach is not our choice of modal base or ordering source, but rather how we induce the ordering based on these parameters. To model the deliberating process of the agent with limited knowledge, we "flatten" the ordering induced so that it does not distinguish between worlds in one and the same cell of the decision problem. In other words, all worlds in which a particular action from the decision problem is taken are ranked together. We obtain this by defining an order $\leq_{g(w)}^{[\delta|f](w)}$ which compares single worlds u and v only according to ordering source propositions that hold throughout their respective cells in the partition induced by δ.

Definition 6 (Deliberative preference ranking)

[30] Kolodny and MacFarlane (2010) define ***chooseable*** as follows (p. 20):

Definition 4 (Chooseable (preliminary – Kolodny and MacFarlane))
ϕ is choosable relative to $\langle w, i\rangle$ iff there is some action specification Δ such that '$\Diamond_e$(Δ is done by agents who know they are doing Δ)' and '$\Box_e$(Δ is done $\supset$ ϕ)' are both true at $\langle w, i\rangle$.

Note that this is 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 circumstantially possible, i.e. '$\Diamond_c\varphi$' at w iff $\neg\Box\neg\varphi$ at $\langle f_c, w\rangle$).

Definition 5 (Chooseable)
ϕ is choosable relative to $\langle w, i\rangle$ iff there is some action specification Δ such '$\Box_e$($\Diamond_c$(Δ is done by agents who know they are doing Δ))' and '$\Box_e$(Δ is done $\supset$ ϕ)' are both true at $\langle w, i\rangle$.

[31] Once the modified semantics for the deliberative modal is in place, we will see that using f_c instead gives rise to exactly the same truth conditions as a circumstantial modal base on the standard application of the Kratzer-style framework as discussed in Section 2.3.1 above.

Let π, A be arbitrary sets of propositions, where the members of π are mutually exclusive. For each world $u \in \bigcup \pi$, let $[u]_\pi$ be the cell in π that contains u. The deliberative preference ranking *induced by A relative to π is a relation $\leq_A^\pi \in (\bigcup \pi)^2$ defined as follows: For all $u, v \in \bigcup \pi$,*

$$u \leq_A^\pi v \text{ iff } \left\{p \in A \middle| [v]_\pi \subseteq p\right\} \subseteq \left\{p \in A \middle| [u]_\pi \subseteq p\right\}$$

As in the case of Kratzer's $\leq_{g(w)}$, it is easy to see that $\leq_A^\pi$ is a pre-order on the worlds in its domain.

Fact 2
For all π, A as in Def. 6, $\leq_A^\pi$ is reflexive and transitive.

To characterize the relevant ranking in our case, we rely on the given deontic ordering source g and fill in the salient partition. Recall that the latter is jointly determined by the decision problem and the modal base. The induced ranking is as follows.

$$(21) \qquad u \leq_{g(w)}^{[\delta|f](w)} v \text{ iff } \left\{p \in g(w) \middle| [v]_{[\delta|f](w)} \subseteq p\right\} \subseteq \left\{p \in g(w) \middle| [u]_{[\delta|f](w)} \subseteq p\right\}$$

We will shortly illustrate how this ordering operates in the miners case. For now, let us point out two main differences between the ranking in (21) and the traditional Kratzer-style ranking $\leq_{g(w)}$. First, Kratzer defines $\leq_{g(w)}$ as a relation in W^2, the set of all pairs of worlds, hence for instance $u \leq_{g(w)} u$ for all $u \in W$ due to the reflexivity. In contrast, our relation $\leq_{g(w)}^{[\delta|f](w)}$ is restricted to the worlds in the modal base $\bigcap f_e(w)$. This is a technicality which could easily be defined away; but since none of the reasoning we are interested in relies on worlds outside of this set, it is innocuous for our purposes.

The second difference between the two rankings arises on the domain for which both are defined, and this difference does play a crucial semantic role. In determining whether $\leq_{g(w)}^{[\delta|f](w)}$ holds between two worlds u, v, not all propositions in $g(w)$ are relevant, but only those which contain their respective cells of the filtered decision problem. To give some intuitive grounding to this idea, consider some preference/value/goal G in $g(w)$: For G to affect the ordering of worlds, there must be some available course of action that (given the agent's information) will secure G. In principle, this is compatible with a number of different ways of comparing actions.[32] In the case of the miners problem, with the particular ordering source we provided, the effect is that in comparing actions in the filtered decision problem with each other, each is only as good as the worst outcome it can possibly have.

[32] One might, for example, consider an approach in terms of expected utilities. This is irrelevant to the problem under discussion, as we assume a 50% chance of the miners being in either of the shafts, and each life is considered equally precious.

Table 5: Parameters for the deliberative account and uncertain information

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

On its deliberative reading, *ought* expresses human necessity with respect to modal base f, ordering source g, and decision problem δ. The truth conditions are analogous to the ones given before, but now the additional parameter $[\delta|f](w)$ is taken into account as a 'filter' when the ordering source induces the pre-order that determines which worlds are best. (22) spells out the adaptation of Kratzer's formulation.[33]

(22) $\boxplus\varphi$ is true at $\langle w, f, g, \delta\rangle$ iff
$\forall v \in \bigcap f(w)\left[\exists u \in \bigcap f(w)\left[u \leq_{g(w)}^{[\delta|f](w)} v \wedge \forall z \in \bigcap f(w)\left[z \leq_{g(w)}^{[\delta|f](w)} u \rightarrow z \in \phi\right]\right]\right]$

With the Limit Assumption we can equivalently define the set of "deliberatively best" worlds as in (23a) and give the semantics of deliberative *ought* as in (23b). Notice that as long as $\leq_{g(w)}$ is finite, there is guaranteed to be a set of "best" worlds, even if the decision problem is not finite.

(23) a. $O(f, g, w, \delta) = \left\{u \in \bigcap f(w) \middle| \forall v \in \bigcap f(w)\left[v \leq_{g(w)}^{[\delta|f](w)} u \rightarrow u \leq_{g(w)}^{[\delta|f](w)} v\right]\right\}$

b. $\boxplus\varphi$ is true at $\langle w, f, g, \delta\rangle$ iff $O(w, f, g, \delta) \subseteq \varphi$.

3.2 Predictions

In the miners scenario, we get the expected predictions. Recall that the ordering source is defined as follows:

[33] Since we assumed that the relevant modal base specifies that the agent gets to act, the decision problem is a partition of the modal background and $\bigcap f_e(w) = \bigcup [\delta|f](w)$. If the modal base did not rule out worlds at which the agent is prevented from acting, the restriction on the domain of quantification in (22) would have to be to the possibly smaller set $\bigcup [\delta|f](w)$.

Table 6: Parameters for the deliberative account and updated information

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

(8) $g(w)$ = {all miners are saved, at least 9 miners are saved, …, at least 1 miner is saved }

Under uncertainty, the set of worlds in which we block shaft A contains both worlds in which the miners are in shaft A (and we save all of them) and worlds in which the miners are in shaft B (and we lose all of them). Therefore, the worlds corresponding to blocking A while the miners are in A are ranked only as high as the unfortunate worlds in which we lose all the miners. In particular, *all miners are saved* does not "count" in favor of these worlds even if we *do* end up saving all ten miners. Similarly for blocking shaft B. Both these cells verify none of the ordering source propositions. In contrast, at all worlds in the cell corresponding to *we block neither shaft* nine miners are saved (be they in A or in B). Therefore, this cell is strictly better than the other two, and a world w can only be strictly better than a world v if it is a blocking-neither world.

Things would change if we learned where the miners are: For example, if the information state is updated with the information that they are in A, this becomes part of $f_e(w)$. Consequently, the cells corresponding to the three actions all contain only worlds in which the miners are in shaft A. If we rank the worlds in this narrower modal background, at all worlds in the cell corresponding to blocking shaft A we save 10 miners, at all worlds in the cell corresponding to blocking neither shaft we save only 9 miners, and at all worlds in the cell corresponding to blocking shaft B we lose all the miners. Consequently, the worlds that are ranked as strictly best are those at which we block shaft A. These predictions are what we wanted.

Our analysis shares an important feature with Kolodny and MacFarlane's account: both the ranking of worlds and the designation of optimal worlds turn out to be ***seriously information dependent*** in their sense. Our analysis however provides a much 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 our view there is a sense in which this is correct and a sense in which it is not correct. In particular, while it is true that the ranking and the optimal worlds vary with information, the ordering source does not. 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}, while if we know that the miners are in A, it is simply {AA}).

3.3 Stability and circumstantial modal bases

Based on the intuition that the truth value of the deliberative reading depends upon the information available to the agent, we have used f_e as the modal base. Now that the analysis is in place, we can easily convince ourselves that by using the circumstantial conversational background f_c instead, we would fail to predict the intuitions about deliberative modals. In particular, if the modal base is set to f_c, our three-parameter version of the Kratzer-style account yields the same prediction as the original two-parameter version (i.e., the one that does not take into account a decision problem). As before, with a circumstantial modal base, which sentences are predicted to be true depends on where the miners are. Assume that they are in shaft A at a world w_A. In that case:

- the circumstantial modal background is $\bigcap f_c(w_A)$ = AA ∪ AB ∪ AN.
- the filtered problem is: $[\delta|f](w)$ = {AA, AB, AN}
- and clearly, AA $<^{[\delta|f](w)}_{g(w)}$ AN $<^{[\delta|f](w)}_{g(w)}$ AB.

This predicts that what we ought to do is block shaft A (and not neither shaft). Analogously, for a world in which they are in shaft B it is predicted that we ought to block shaft B.

We think that this insensitivity of the circumstantial modal base to the decision problem should follow in general from assumptions about the circumstantial conversational background. But there is a certain vagueness in the standard Kratzerian specification of the circumstantial conversational background as the function which assings to each world w a set of propositions holding at w that specify the "relevant circumstances." While it may not always be clear what circumstances are relevant (independently from considering what is needed to derive speaker intuitions about modalized sentences), it seems to us that at least all those facts should

count as relevant that are needed to evaluate the contextually given decision problem according to one's criteria (preferences, goals, rules, in short, the value of the ordering source). This means that the circumstantial modal base should have the following property, which we call ***persistence***.

Definition 7 (Persistence)
A modal base f is persistent for a decision problem δ w.r.t. an ordering source g at a world w iff for any proposition $p \subseteq W$ and worlds $u, v \in \bigcap [f+p](w)$, $u \leq_{g(w)}^{[\delta|f+p](w)} v$ iff $u \leq_{g(w)}^{[\delta|f](w)} v$.

One way in which persistence may be ensured is for all propositions in the ordering source to come into play in comparing worlds. This is the case whenever none of them "cuts across" cells in the partition any longer. We state this fact without proof.[34]

Fact 3
For all f, δ, g, w, it holds that f is persistent for δ with respect to g if for all $v \in \bigcap f(w)$ and $p \in g(w)$, if $[v]_{[\delta|f](w)} \cap p \neq \emptyset$ then $[v]_{[\delta|f](w)} \subseteq p$.

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 at w with respect to g and δ, this means that the judgements for deliberative *ought* with respect to this parameter collapse with the ones for two-parameter ought with the same g and the circumstantial modal base, which we considered the appropriate rendering for objective *ought*.

4 Relation to the standard Kratzer analysis

Kratzer's framework for the semantics of modal expressions in natural language relies on two conversational backgrounds, namely a modal base and an ordering source. In order to extend her analysis to deliberative *ought* as evidenced by (2a) in the miners scenario, we added a decision problem as a third parameter. Together with the other two parameters, the decision problem is used to induce a pre-order on the worlds in the modal background which can, but need not, agree with the one induced by the ordering source alone. Thus it may seem that we have achieved an account of deliberative modality by departing considerably from the standard picture. In the following, we want to reflect on the exact extent of our departure from

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

the standard version of Kratzer's framework. This question strikes us as important as it is generally considered an appealing factor of the system that it avoids ambiguities while offering a powerful system to account for all kinds of modal expressions in natural language. In particular, we want to evaluate two issues: First, does the third parameter force us to adopt ambiguities in order not to affect the interpretation of non-deliberative *ought*? Second, does appealing to a specific parameter in evaluating deliberative *ought* undermine claims about the unity and interrelation of the modal system? We think that the answer to both questions is negative.[35] To show this, we want to highlight two ways in which our analysis and the standard picture can be fully reconciled.

4.1 Decision problems

Clearly, not all occurrences of *ought* appear to be deliberative, both in light of our discussion and on independent grounds. In the context of the miners scenario, we individuated an objective reading which gave rise to different intuitions about the truth-conditions of (2a) and (4) in the miners scenario than are observed on the deliberative reading. Others (Schroeder, 2010), have defended ***ambiguity*** theories for *ought* that introduce a lexical distinction between deliberative *ought* as in examples like (2a), and non-deliberative *ought* as in (24):

(24) The car ought to start in less than five minutes.

It is irrelevant to our current issue whether this sentence has a genuine deontic interpretation, or if the modal is truly epistemic (in a sense that would require a doxastic, stereotypical etc. ordering source in a Kratzer-style analysis). We want to know if the decision problem has a role to play in such a case, or if a different analysis is required (for example, two and three parameterized) lexical entries for *ought*.

In one sense, the answer to the question whether "the decision problem" plays any role in the interpretation of (24) must be negative. After all, cars don't make decisions. The more interesting question is whether the sort of coarsening of logical space that is introduced by the technical parameter thus labeled is useful (or even necessary) in analyzing this sort of interpretation of *ought*. But even that question we think we can remain entirely neutral on. We motivate this stance by showing that, even if the decision problem turns out to be entirely irrelevant, there is a natural setting of the parameter on which our semantics collapses on Kratzer's,

[35] Note that, regarding the first question, we do not want to commit us to whether ambiguities may need to be adopted for other reasons. The scope of our claim is just that the possibility of introducing the third parameter does not in itself require the introduction of ambiguities.

independently of how modal base and ordering source are chosen.

The decision problem parameter is blended out if δ takes a "default" non-agential value:[36,37] Let δ^* be the trivial partition of the set of possible worlds each of whose cells contains exactly one world (i.e., δ^* is the finest possible partition of the logical space). If we adopt δ^* as a default setting for the decision problem parameter in all cases in which Kratzer's two parameter semantics gave the right result, our framework can be seen as a generalization of Kratzer's apparatus. Then both the standard setting (with δ^*) and our application to deliberative *ought* (with a non-trivial value for the decision problem parameter) come out as applications of one and the same three parameterized theory.

It is easy to prove that for worlds in the modal background, the ordering induced by δ^* relative to parameters f, g, w just is the same as the ordering induced by f, g, w according to the standard Kratzer recipe. Let us first define the maximally specific decision problem.

Definition 8 (Maximally specific decision problem)
The maximally specific *decision problem is the partition* $\delta^* = \{\{w\} \mid w \in W\}$.

The claim here is that, given any modal base f, ordering source g and world of evaluation w, the Kratzerian ordering $\leq_{g(w)}$ matches the ordering $\leq_{g(w)}^{[\delta^*|f](w)}$. The two orderings are not exactly identical because $\leq_{g(w)}^{[\delta^*|f](w)}$ differs on worlds outside of the modal background, but they must agree on all worlds inside of it.

Theorem 2
For all $u, v \in \bigcap f(w)$, $u \leq_{g(w)} v$ *iff* $u \leq_{g(w)}^{[\delta^*|f](w)} v$.

Proof. Via the standard definition of $\leq_{g(w)}$ we derive:

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

□

[36] Recall from Section 3.3 above that a collapse in truth-conditions is also obtained if the modal base provides enough information such that all ordering source propositions are considered in the comparison of any two worlds in the modal background.

[37] Note that this trivial decision problem must be exempt from the requirement of "choosability" discussed in Footnote 30. We assume that only choosable decision problems can be contextually salient under ordinary circumstance.

The significance of this result is that we need not take deliberative *ought* to have its own distinctive lexical entry, even if it turns out that other interpretations of *ought* do not make direct use of the decision problem. We can simply assume that, in those cases, the decision problem is set to δ^*.[38] In principle, nothing prevents even other modal operators from also being sensitive to decision problems as long as – if they do not possess deliberative readings – the parameter is specified lexically to δ^*. 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 Trimming

Above, we showed that Kratzer's standard two-parameter semantics can easily be derived from a general three-parameter framework with a particular default setting for the decision problem. In this section we want to turn things around and ask if there is a way of deriving our result in the standard two-parameter setting. Note that we have already proved that this is impossible if we want the theory to reflect certain intuitions, in particular, that, in the miners scenario, learning where the miners are would change the agents' judgments as to what is the best course of action. We already know that such a theory cannot be obtained in view of the result that Kratzer's semantics is not seriously information dependent.

Nevertheless, putting aside these intuitions about the roles played by the modal base (as reflecting information) and the ordering source (as reflecting goals, preferences, or rules), it is obvious that our three parameters together serve to induce a pre-order on the modal base. In contrast, on the standard two-parameter version of the Kratzer framework, a pre-order on the set of possible worlds was induced by the ordering source alone. Couldn't we find an ordering source that derives the desired pre-order as obtained from the interaction of the three parameters?

Recall that the usual values for the modal base are functions specifying what is known to some agent (epistemic) or what the relevant circumstances are (circumstantial), whereas the ordering source is supplied as a function specifying what is commanded (deontic), what someone wants (bouletic), etc. Since the decision affects the derived pre-order, if we are to omit it, we would need to simulate its effect

[38] Recall that we made a similar move above, in Section 2.1, treating the necessity operator $\Box$ as a special case of the human necessity operator $\boxplus$, obtained by setting the ordering source to the trivial conversational background $\lambda v.\emptyset$.

by modifying the remaining parameters. It is indeed possible to achieve this by using a "trimmed" variant of the original ordering source.

Definition 9 (Trimmed ordering source)
Given a set π of mutually exclusive propositions, for any proposition p, let the corresponding π-trimmed proposition $p^{\downarrow\pi}$ *be defined as follows:*

$$p^{\downarrow\pi} = \{w \in W | [w]_\pi \subseteq p\} \tag{i}$$

Furthermore, for any set A of propositions, the corresponding π-trimmed set is defined as

$$A^{\downarrow\pi} = \{p^{\downarrow\pi} | p \in A\} \tag{ii}$$

Finally, for any ordering source g and filtered decision problem $[\delta|f]$, the $[\delta|f]$ -trimmed ordering source *is a function $g^{\downarrow[\delta|f]}$ such that for all $w \in W$,*

$$g^{\downarrow[\delta|f]}(w) = g(w)^{\downarrow[\delta|f](w)} \tag{iii}$$

The trimmed ordering source is a conversational background that can be described as "our preferences as compared to the information we have on our choices." Informally, it is obtained by "cutting away" 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. Like any other conversational background, $g^{\downarrow[\delta|f]}(w)$ can be used to induce a preorder $\leq_{g^{\downarrow[\delta|f]}(w)}$ on worlds in the usual way. Instead of relativizing to the decision problem whenever we compare two worlds, we can also globally use $g^{\downarrow[\delta|f]}(w)$ as the ordering source for *ought* on 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 we are envisaging. Arrows represent determination relations, the label on nodes refer to our definitions of those relations.

It can be shown that the two approaches are equivalent for the evaluation of any modal operator that is not modified by a conditional antecedent. That is, if $g^{\downarrow[\delta|f]}(w)$ constitutes the contextually salient ordering source, the standard two-parameter Kratzer semantics yields exactly the same predictions as our version with the explicit additional parameter.

Theorem 3
For any worlds w, u, v and conversational backgrounds δ, f, g s.t. $u, v \in \bigcap f(w)$:

$$u \leq^{[\delta|f](w)}_{g(w)} v \text{ iff } u \leq_{g^{\downarrow[\delta|f]}(w)} v$$

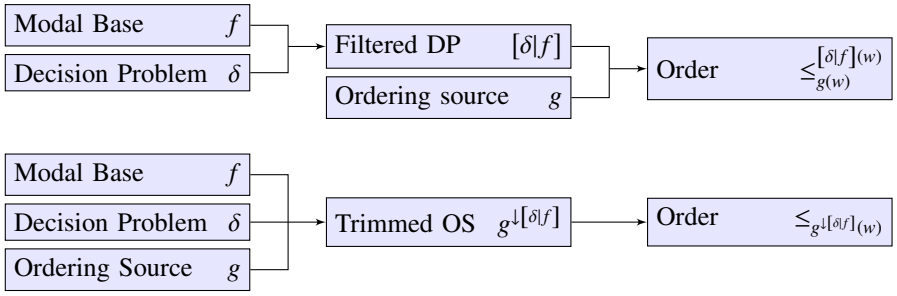

Figure 1: Architecture of our account of deliberative modality (top) and the alternative using a trimmed ordering source (bottom).

Proof.

$$
\begin{aligned}
u \leq_{g(w)}^{[\delta|f](w)} v \text{ iff } & \left\{ p \in g(w) \,\middle|\, [v]_{[\delta|f](w)} \subseteq p \right\} \subseteq \left\{ p \in g(w) \,\middle|\, [u]_{[\delta|f](w)} \subseteq p \right\} && \text{(Def. 6)} \\
\text{iff } & \left\{ p \in g(w) \,\middle|\, v \in p^{\downarrow[\delta|f](w)} \right\} \subseteq \left\{ p \in g(w) \,\middle|\, u \in p^{\downarrow[\delta|f](w)} \right\} && \text{(Def. 9i)} \\
\text{iff } & \left\{ p^{\downarrow[\delta|f](w)} \,\middle|\, p \in g(w), v \in p^{\downarrow[\delta|f](w)} \right\} \subseteq \left\{ p^{\downarrow[\delta|f](w)} \,\middle|\, p \in g(w), u \in p^{\downarrow[\delta|f](w)} \right\} \\
\text{iff } & \left\{ q \in g(w)^{\downarrow[\delta|f](w)} \,\middle|\, v \in q \right\} \subseteq \left\{ q \in g(w)^{\downarrow[\delta|f](w)} \,\middle|\, u \in q \right\} && \text{(Def. 9ii)} \\
\text{iff } & \left\{ q \in g^{\downarrow[\delta|f]}(w) \,\middle|\, v \in q \right\} \subseteq \left\{ q \in g^{\downarrow[\delta|f]}(w) \,\middle|\, u \in q \right\} && \text{(Def. 9iii)} \\
\text{iff } & u \leq_{g^{\downarrow[\delta|f]}(w)} v && \square
\end{aligned}
$$

Although we have just shown that the three-parameter version of *ought* can be reduced to the two-parameter version with a somewhat more abstract ordering source, we want to point out a few observations in favor of making explicit the dependence of the ordering on the decision problem.

First, Kratzer's conversational backgrounds are usually required to be contextually salient (consider in particular implementations in DRT where the converational 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 have presented it as one of the main advantages of our framework over Kolodny and MacFarlane's analysis that the preferences that underly the agent's deliberative process play a role transparently. As on the standard two-parameter Kratzer framework, this allows us to systematically predict the interaction between preferences and circumstances. In forcing three-parameter de-

liberative *ought* into the two-parameter mold, we lose this advantage: in particular, learning where the miners are amounts to a change not only in the modal base, but also in the ordering source. In contrast, if the decision problem is introduced separately, the preferences can be kept constant.

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 (OCO construal). In our three-parameter version, as the filtering of the decision problem depends on the modal base, an update of the modal base may effect a the order $\leq_{g(w)}^{[\delta|f](w)}$ 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).[39] In our opinion this would constitute a more dramatic departure from the spririt of Kratzer's approach than our addition of a third parameter.

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 is not seriously information dependent and lacks a clear solution to the miners puzzle. 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

[39] Sæbø (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 Fintel and Iatridou (2005).

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.

References

Belnap, N. and J. Horty. 1995. The deliberative *Stit*: a study in the logic of action, omission, ability and obligation. *Journal of Philosophical Logic*, 24:583–644.

Belnap, N., M. Perloff, and M. Xu. 2001. *Facing the Future*. Oxford University Press.

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*.

DeRose, K. 1991. Epistemic possibilities. *The Philosophical Review*, 100:581–605.

Egan, A. 2007. Epistemic modals, relativism, and assertion. *Philosophical Studies*, 133:1–22.

von Fintel, . and S. Iatridou. 2005. What to do if you want to go to Harlem: Anankastic conditionals and related matters.

von Fintel, S. and A. Gillies. 2007. An opinionated guide to epistemic modality. *Oxford Studies in Epistemology*, 2:32–62.

von Fintel, S. and A. Gillies. 2008. CIA leaks. *Philosophical Review*, 117:77–98.

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.

Hacking, I. 1967. Possibility. *The Philosophical Review*, 76:143–168.

Hamblin, C. 1958. Questions. *Australasian Journal of Philosophy*, 36:159–168.

Horty, J. 2001. *Agency and Deontic Logic*. Oxford University Press.

Jackson, F. 1991. Decision theoretic consequentialism and the nearest dearest objection. *Ethics*, 101(3):461–482.

Kaufmann, S. and M. Schwager. t.a. A unified analysis of conditional imperatives. In *Proceedings of SALT 19*. 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 Philosphy*, 3(1):337–355.

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.

Lewis, D. 1973. *Counterfactuals*. Blackwell, Oxford.

MacFarlane, J. to appear. Epistemic modals are assessment-sensitive. In Weatherson and Egan (to appear).

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.

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*. Akademieverlag, Berlin.

Schroeder, M. 2010. Oughts, agents and actions. *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.

Stephenson, T. 2007. Judge dependence, epistemic modals, and predicates of personal taste. *Linguistics and Philosophy*, 30:487–525.

Weatherson, B. and A. Egan, editors. to appear. *Epistemic Modality*. Oxford University Press, Oxford.

Yalcin, S. 2007. Epistemic modals. *Mind*, 116(4):983–1027.

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*.