What we know and what to do

Nate Charlow

Abstract   This paper discusses an important puzzle about the semantics of indicative conditionals and deontic necessity modals (*should*, *ought*, etc.): the Miner Puzzle (Parfit, ms; Kolodny and MacFarlane, J Philos 107:115–143, 2010). Rejecting modus ponens for the indicative conditional, as others have proposed, seems to solve a version of the puzzle, but is actually *orthogonal* to the puzzle itself. In fact, I prove that the puzzle arises for a variety of sophisticated analyses of the truth-conditions of indicative conditionals. A *comprehensive* solution requires rethinking the relationship between relevant information (what we know) and practical rankings of possibilities and actions (what to do). I argue that (i) relevant information determines whether considerations of value may be treated as reasons for actions that realize them and against actions that don't, (ii) incorporating this normative fact requires a revision of the standard ordering semantics for weak (but not for strong) deontic necessity modals, and (iii) an off-the-shelf semantics for weak deontic necessity modals, due to von Fintel and Iatridou, which distinguishes "basic" and "higher-order" ordering sources, and interprets weak deontic necessity modals relative to both, is well-suited to this task. The prominence of normative considerations in our proposal suggests a more general methodological lesson: formal semantic analysis of natural language modals expressing normative concepts demands that close attention be paid to the nature of the underlying normative phenomena.

Keywords   Conditional obligation · Deontic modality · Weak and strong deontic necessity · Indicative conditionals · Ordering semantics

N. Charlow (✉)
University of Toronto, 170 St. George Street, Toronto, ON M5R 2M8, Canada
e-mail: nate.charlow@gmail.com

⌂ Springer

1 Introduction

This paper discusses a puzzle about indicatives and deontic modals (*should*, *ought*, etc.), the Miner Puzzle, described in a recent paper of Niko Kolodny and John MacFarlane. It is a little puzzle—tidy and easy to summarize—but it has ramifications, of a foundational nature, for theories of both indicative conditionals and deontic modals. Rejecting modus ponens for the indicative conditional, as Kolodny and MacFarlane propose, seems to solve a version of the puzzle, but is actually explanatorily orthogonal to the puzzle itself. I argue that a genuinely explanatory solution to the Miner Puzzle requires rethinking the relationship between relevant information (what we know) and practical rankings of possibilities and actions (what to do). While Kolodny and MacFarlane do endorse a particular conception of this relationship (albeit incidentally), I argue that it does not, in the end, yield a genuinely explanatory solution to the Miner Puzzle.

Not so with the account developed and defended herein. On that account, relevant information is seen to make a twofold difference in thinking about what to do. It restricts the background of salient possibilities against which deliberation occurs. But it also determines whether considerations of value *may be treated as reasons* for actions that realize them and against actions that don't. The Miner Puzzle is not a puzzle about modus ponens. It is, rather, about the special sensitivity of certain ways of expressing obligation in natural language (and, so, certain kinds of obligation) to relevant information—a sensitivity that traces to informational constraints on treating considerations of value as reasons.

In Sect. 2, I describe the puzzle and rule out some first attempts at resolving it. I argue, with Kolodny and MacFarlane, that a solution ultimately requires a genuine rejection of modus ponens for the indicative conditional. But, in Sect. 3, I show that this is actually explanatorily orthogonal to the Miner Puzzle. This is supported by appeal to some results about the semantics of deontic conditionals. The thrust of these results is that a puzzle effectively identical to Kolodny and MacFarlane's will arise, even on state-of-the-art treatments of indicatives (e.g., those of Angelika Kratzer and Thony Gillies) that invalidate modus ponens. In Sect. 4, I develop this point, highlighting three, specific respects in which Kolodny and MacFarlane's resolution of the puzzle is theoretically less-than-satisfactory. This is not to say they are *wrong*, or that their commitment to their account would prevent them from accepting the central elements of mine. It is only to say that the story about the puzzle that they offer will be but a peripheral part of a fully explanatory account.

Finally, in Sects. 5 and 6, I try to tell what I take to be the whole story. On my account, the puzzle really has to do with the special sensitivity of *weak* deontic necessity modals (*should*, *ought*, etc., in contrast to strong modals like *must*, *have to*, etc.) to relevant information. The nature of this sensitivity can, I show, be modeled in a theoretically satisfying, but still precise, way. One promising way to do this is to avail ourselves of a key insight from a linguistically influential treatment of the meaning of weak necessity modals, due to Kai von Fintel and Sabine Iatridou (and, ultimately, an old idea of Aaron Sloman). On the version of this treatment that I develop, what we *should* do is a function of first-order concerns as well as higher-order concerns (concerning which first-order concerns it is appropriate to treat as reasons in favor of

actions that realize them). This idea dovetails nicely with some recent normative work on the nature of the relationship between knowledge and practical reasons.

Although this paper's ostensible subject-matter is the *semantics* of indicatives and deontic modals, a variety of *normative* and *decision-theoretic* considerations end up figuring prominently in our discussion. While this may seem strange, such considerations play an important role in motivating the semantic analysis we go on to propose. I take there to be a more general methodological moral here. Satisfactory formal semantic analysis of natural language modals expressing normative concepts sometimes demands a more sophisticated understanding of the nature of the underlying normative phenomena.

2 Ground rules

In "Ifs and Oughts", Kolodny and MacFarlane (hereafter 'K&M') identify a pressing puzzle about indicative conditionals and deontic necessity operators (*should*, *ought*, etc).[1] The setup: ten miners are trapped in a shaft—A or B, although we do not know which—and threatened by rising waters. We can block one shaft or neither, but not both. If we block the correct shaft, everyone lives. If we block the wrong shaft, everyone dies. If we do nothing, only one miner dies. The decision matrix is as follows (with *in_X* abbreviating the proposition *the miners are in X*, *block_X* the proposition *we block X*).

	in_A	*in_B*
block_A	All live	All die
block_B	All die	All live
¬(*block_A* ∨ *block_B*)	Nine live	Nine live

Let *c* be a context (of utterance or assessment, it makes no difference for my purposes) where we're discussing what to do about the miners, □φ abbreviate *we should see to it that φ*, and (*if φ*)(*ψ*) abbreviate an indicative conditional whose antecedent is φ and consequent is ψ. There is a strong intuition, reliably corroborated by informants, that the sentences in (1) are all true at *c* (so, a fortiori, *consistent* at *c*).

(1)   a.   *in_A* ∨ *in_B* (They're either in A or B)
        b.   (*if in_A*)(□*block_A*) (If they're in A, we should block A)
        c.   (*if in_B*)(□*block_B*) (If they're in B, we should block B)
        d.   ¬(□*block_A* ∨ □*block_B*) (We shouldn't block either)

But there is a *proof* that seems to show this can't be right.

Proof Suppose that *in_A* ∨ *in_B*, (*if in_A*)(□*block_A*), and (*if in_B*)(□*block_B*). Now suppose *in_A*. Then, by modus ponens, □*block_A*. Then, by ∨- INTRO, □*block_A* ∨ □*block_B*. Now suppose *in_B*. Then, by modus ponens, □*block_B*. Then, by ∨- INTRO, □*block_A* ∨ □*block_B*. So, by ∨- ELIM, □*block_A* ∨ □*block_B*. □

---

[1] K&M credit the example to Parfit (ms), but the presentation of the puzzle as one about indicatives and *ought* is, I believe, original to them. Their presentation of the puzzle utilizes the weak modal *ought*, while mine utilizes the weak modal *should*. Nothing turns on this; I just find *should* more natural.

So we have a conflict between a firm, and widely shared, intuition of truth (hence, of consistency) and an apparent proof of inconsistency. That, basically, is the puzzle. There are three kinds of response one could have to the puzzle.

CHALLENGING (1d)
Rejecting that neither shaft should be blocked (while accepting the proof's soundness and validity).

CHALLENGING SOUNDNESS
Rejecting ≥1 of the proof's premises (while accepting its validity).

CHALLENGING VALIDITY
Rejecting one of the proof's inferences (while accepting each of its premises).

What sorts of considerations would motivate each kind of response? Challenging (1d), K&M suggest, is motivated by an *Objectivist* view of *should* (p. 117):

OBJECTIVE SHOULD
S should Φ iff, roughly, Φ-ing is desirable from a vantage of full information with respect to the relevant facts.

In the miner case, the relevant facts are represented as the columns of the decision matrix; full information with respect to these facts means knowing the miners' location, in which case one or the other shaft (the one housing the miners) will be such that blocking it is desirable. K&M make short work of this response: "[I]n deciding what we ought to do, we always have limited information… Thus the objectivist's [*should*] seems useless in deliberation" (K&M, p. 117). Although Objectivism has a venerable pedigree, this critique, sufficiently spelled out, seems to me persuasive. I will be setting aside Objectivism for the rest of the paper.

Challenging the proof's soundness, while accepting its validity, means rejecting one of the conditional obligation claims, (1b) or (1c). This, they suggest, is motivated by a *Subject-Sensitive Subjectivism* for *should* (K&M, p. 118):

SUBJECT- SENSITIVE SUBJECTIVE SHOULD
S should Φ iff, roughly, Φ-ing is desirable from the vantage of *S*'s information.

For the Subject-Sensitive Subjectivist, either (1b) or (1c) must be false. For suppose otherwise. The miners are either in A or B. Suppose the former. Then, by modus ponens, since (1b) is true, we should block A. And the same goes if they are in B and (1c) is true. But, since blocking neither is preferred (from the vantage of our limited information), the Subjectivist claims this is the course of action we should choose. On Subject-Sensitive Subjectivism, our supposition generates a contradiction. So either (1b) or (1c) must be false.

K&M argue against Subject-Sensitive Subjectivism, on the grounds that it fails to make sense of the phenomenon of disagreement from a position of epistemic advantage: someone who knows the miners are in B can felicitously disagree with someone who claims we shouldn't block either, by uttering (2) (K&M, p. 119).

(2)   Actually, you should block B.

This is the case even though, by the Subject-Sensitive Subjectivist's lights, she is saying something she *knows to be false*.

But raw disagreement "data" (of the sort based on intuitions about whether disagreement of a specific type is present or felicitous) is extremely fraught,[2] and it is a good general policy not to rest too much on it. So let me add that there is clearly a reading of the indicatives—call it the "Ramsey reading"—on which they are both *obviously* acceptable. If you have trouble getting such a reading, here is some prodding. The Ramsey test for the acceptability of an indicative "invites us to add the information carried by the antecedent to the contextually relevant stock of information… and check the fate of the consequent" (Gillies 2010, p. 27). On the Ramsey reading of (1b), we add the information carried by its antecedent to the relevant stock of information, and subsequently check whether the consequent of (1b) is, from the vantage of this enriched body of information, true. And, of course, it is: if our information settles the miners' location, it also settles which shaft we should block. So, the Ramsey readings of (1b) and (1c) are clearly true. Although Subject-Sensitive Subjectivism fails to capture these readings, they are *no less real for that*; indeed, given the wide convergence of intuitions on the consistency of the sentences in (1), the Ramsey readings of (1b) and (1c) seem rather more prominent than the Subject-Sensitive Subjectivist's claimed readings. So Subject-Sensitive Subjectivism fails to defuse the puzzle of explaining the apparent consistency of the sentences in (1).[3]

This leaves us with challenges to the proof's validity. Since ∨- INTRO and ∨- ELIM are truth-preserving (on their classical, truth-functional interpretations), this leaves only one option: denying modus ponens for some natural language indicatives of the form ⌜(*if* ϕ)(ψ)⌝. Such a denial can take any of the following forms:

WIDE- SCOPING
Some indicatives—in particular, some whose consequents contain modals—do not receive a standard semantics. Rather, they are assigned logical forms (LFs) in which their modals take scope over a material conditional.

MODALS AS RESTRICTABLE QUANTIFIERS
Indicatives are *uniformly* mapped to non-standard LFs. Specifically, indicative antecedents are interpreted as domain-restrictors for a quantifier Q that does not validate detachment (i.e., does not validate the inference form $Qx(Ax)(Bx)$, Ax / Bx).

---

[2] For critiques of MacFarlane's use of such data, see von Fintel and Gillies (2008), Dreier (2009).

[3] There is another line of response, which holds that the proof's conclusion is compatible with (1d), since there is, at *c*, a true Subject-Sensitive Subjectivist reading of (1d) and a true Objectivist reading of the proof's conclusion. This line of response is compatible with thinking that *should* is lexically ambiguous between Subjectivist and Objectivist meanings, or with thinking that *should* has a systematically context-dependent meaning, such that, in some contexts, it has the Subjectivist meaning, and in other contexts the Objectivist meaning (cf. Kratzer 1981). As K&M (Sect. II) argue, neither tack ameliorates the problems for Subjectivism or Objectivism.

NARROW- SCOPING

Indicatives *uniformly* retain their standard LFs (and modals in consequent position remain so at LF). The canonical Stalnaker-Lewis semantics for indicatives, which does validate modus ponens, is revised so that it fails to validate it.

We will dispatch of the wide-scoping resolution here, and go on to consider the other two ways of denying modus ponens for indicatives in more detail below.

Wide-scoping is a *non-radical* way of denying modus ponens. Wide-scopers split natural language indicatives into two disjoint classes: those that receive *genuinely conditional* LFs, and those which do not. Wide-scoping is perfectly compatible with accepting modus ponens for genuinely conditional indicatives (i.e., indicatives which receive genuinely conditional LFs). It simply claims that some natural language indicatives are not genuinely conditional; they are, rather, logically *modal*.

Wide-scoping does block the proof: from $\Box(in\_A \supset block\_A)$ and in_A, it of course does not follow, in normal treatments of deontic modal logic, that $\Box block\_A$. But it will be useful to see, from a semantic perspective, why this is so.

It is standardly held that deontic modals are "doubly relative," in that they are interpreted with respect to a modal base and some sort of selection function, both supplied by a context (see, e.g., Kratzer 1981). Specifically, deontic modals are semantically *quantifiers*, whose domain is generated by applying the selection function to the modal base at a point of evaluation. We'll assume, *very* roughly, that:

– The modal base provides the set of possibilities that are relevant at a context and world (roughly, a body of information considered relevant at the context).[4]
– The selection function selects the deontically best relevant possibilities.

Formally:

– A *context c* determines a pair $\langle F_c, D_c \rangle$
  $F_c = \lambda i. \{ j : j$ is a relevant possibility at $\langle c, i \rangle \}$ is a *modal base* (a function from a possibility i into a set of worlds, namely, the set of i-relevant possibilities/the i-relevant information)
  D_c is a *selection function* choosing the deontically best indices from a domain

The selection function, unlike the modal base, is assumed to be world-independent; this simplifies things somewhat, and nothing turns on it. We will assume, as is also standard:

REFLEXIVITY OF RELEVANCE
$i \in F_c(i)$ (the actual world is always relevant)

---

[4] Again, whether they are contexts of evaluation or assessment is immaterial. Now is a good time to note that everything I say in this paper is officially neutral about the correctness of:

● Relativist accounts, which treat relevant possibilities as given by a parameter of the index: a judge (Stephenson 2008) or information state (MacFarlane 2011; Yalcin 2007)
● Shifty accounts, which treat indicative antecedents as, roughly, Kaplanian monsters—devices for shifting the context of interpretation for a non-restrictable modal (Gillies 2010)

The formal properties of relativist accounts are in all relevant respects similar to contextualist accounts. The predictions of shifty accounts are equivalent, in the relevant cases, to those of domain-restrictor accounts (on their differences, see Gillies 2010). So, while our own solution will be formulated within a shifty account of indicatives, it could easily be formulated as a domain-restrictor account.

REALISM

$D_c(F_c(i)) \subseteq F_c(i)$ (the best worlds are always relevant)

DEFINEDNESS

$D_c(F_c(i))$ is defined and non-empty[5]

Deontic necessity modals are assumed, as is standard, to be universal quantifiers over the domain defined by the modal base and selection function:

Definition 1 A modal formula of the form $\ulcorner \Box \phi \urcorner$ is true at a context c and world i (notation: $[\![\Box \phi]\!]^{c,i} = 1$) iff $D_c(F_c(i)) \subseteq [\![\phi]\!]^c$. (Note: $[\![\phi]\!]^c$ gives the set of worlds j such that $[\![\phi]\!]^{c,j} = 1$.)

Notice that modus ponens is truth-preserving for natural language indicatives only if, for any indicative of the form $\ulcorner (if \ \phi)(\psi) \urcorner$, $[\![\phi]\!]^{c,i} = [\![(if \ \phi)(\psi)]\!]^{c,i} = 1$ implies $[\![\psi]\!]^{c,i} = 1$. But if the LF of an indicative $(if \ \phi)(\Box \psi)$ is in fact $\Box(\phi \supset \psi)$, the fact that the indicative and its antecedent are both true implies, on this semantics, nothing whatsoever about whether $\Box \psi$ is true.

Proof Suppose $D_c(F_c(i)) \subseteq [\![\neg \phi]\!]^c$ (i.e., all the best worlds are $\neg \phi$-worlds), that $[\![\phi]\!]^{c,i} = 1$, and that $\neg \phi$ implies $\neg \psi$. Then $D_c(F_c(i)) \subseteq [\![\phi \supset \psi]\!]^c$, hence, by Definition 1, $[\![\Box(\phi \supset \psi)]\!]^{c,i} = 1$. But since $\neg \phi$ implies $\neg \psi$ (i.e., $[\![\neg \phi]\!]^c \subseteq [\![\neg \psi]\!]^c$) and $D_c(F_c(i)) \subseteq [\![\neg \phi]\!]^c$, $D_c(F_c(i)) \subseteq [\![\neg \psi]\!]^c$. So $[\![\Box \psi]\!]^{c,i} = 0$. Indeed, $[\![\Box \neg \psi]\!]^{c,i} = 1$. $\square$

So giving (1b) and (1c) wide scope LFs blocks the proof. But that is small comfort, as the proposal still predicts the sentences in (3) inconsistent, if we assume (as is, in fact, the case) that all the c-relevant worlds are worlds where the miners are in one of the shafts (i.e., that $F_c(i) \subseteq [\![in\_A \vee in\_B]\!]^c$) (cf. K&M, p. 124).

(3)   a. $(if \ in\_A)(\Box block\_A)$ (If they're in A, we should block A)
      b. $(if \ in\_B)(\Box block\_B)$ (If they're in B, we should block B)
      c. $\neg \Box(block\_A \vee block\_B)$ (Not: we should block at least one)

Proof Suppose that the sentences in (3) are all true, and that (3a) and (3b) have wide scope LFs. Then, by Definition 1, $D_c(F_c(i)) \subseteq [\![in\_A \supset block\_A]\!]^c$ and $D_c(F_c(i)) \subseteq [\![in\_B \supset block\_B]\!]^c$. By Realism, since $F_c(i) \subseteq [\![in\_A \vee in\_B]\!]^c$, $D_c(F_c(i)) \subseteq [\![in\_A \vee in\_B]\!]^c$. Whence $D_c(F_c(i)) \subseteq [\![block\_A \vee block\_B]\!]^c$. Contradiction. $\square$

If you think that (1d) is true in the miner case, then, of course, you should think that (3c) is true too. The wide-scoper's prediction here is maybe a bit less jarring than the prediction that the sentences in (1) are provably inconsistent, but not by much.

3 Denying modus ponens

The upshot, so far, is that handling the puzzle requires a genuine denial of modus ponens—not of the non-radical sort associated with wide-scoping, but a denial of

---

[5] Definedness is the Limit Assumption (Lewis 1973). Rejecting it affects none of my arguments, and would complicate the semantics. See Kratzer (1981), Swanson (2008) for discussion.

modus ponens for *genuinely conditional indicatives* (i.e., ones with conditional LFs). In this section, I canvass the two ways of doing this mentioned in the prior section. Although denying modus ponens is, of course, striking and worthy of the attention which K&M lavish on it, I will show here that it is actually explanatorily orthogonal to the Minor Puzzle (a point that, we will see, is in fact implicitly recognized by K&M). I will show this by proving a few results about the semantics of deontic conditionals. Together, these results show that a puzzle effectively identical to K&M's still arises on state-of-the-art treatments of indicatives that invalidate modus ponens.

A note to the reader: this section is technical. You will not lose terribly much by skipping ahead to Sect. 4. That is, provided you are prepared to trust me that, given a plausible constraint on the deontic selection function (which goes by the name "Stability," and which is introduced in Sect. 3.1), the sentences in (3) are provably inconsistent, even on state-of-the-art treatments of indicatives that invalidate modus ponens.

3.1 Restrictor semantics

Angelika Kratzer famously presented linguistic evidence that "*If*-clauses are devices for *restricting the domains of [quantificational, e.g., modal] operators*" (Kratzer 1991, p. 656).[6] We will call Kratzer's treatment of modals a *generalized quantifier* treatment (since she analyzes modals as generalized quantifiers, taking both restrictor and scope arguments), and her treatment of indicative antecedents a *Restrictor* treatment (since they fill the restriction argument for a generalized quantifier). Together, they comprise a *Restrictor Semantics* for modalized indicatives.

The Restrictor Semantics for natural language indicatives of the form $\ulcorner (if\ \phi)(\Box \psi) \urcorner$ interprets $\Box$ as a restrictable universal quantifier over deontically best indices in the modal base, and indicative antecedents as restrictors of this quantifier's domain.

Definition 2 $[\![(if\ \phi)(\Box \psi)]\!]^{c,i} = 1$ iff $D_c(F_c(i) \cap [\![\phi]\!]^c) \subseteq [\![\psi]\!]^c$

Informally, Definition 2 says that $(if\ \phi)(\Box \psi)$ is true iff all the best ϕ-worlds are ψ-worlds. At LF, $(if\ \phi)(\Box \psi)$ is represented as a *binary modal formula*, something like $\Box(\phi)(\psi)$, where ϕ specifies a domain restriction on $\Box$, and ψ specifies the condition asserted for $\Box$-many individuals in the resulting domain. Unembedded ("bare") modal sentences are treated as vacuously restricted, so that $[\![\Box \psi]\!] := [\![(if\ \top)(\Box \psi)]\!]$. The relationship between bare and embedded modals is, therefore, analogous to that between *everyone loves their mother* and *every person under 12 loves their mother*.

The Restrictor Semantics clearly invalidates modus ponens for indicatives.

Proof Suppose $D_c(F_c(i)) \subseteq [\![\neg\phi]\!]^c$ (the best worlds are $\neg\phi$-worlds), $D_c(F_c(i) \cap [\![\phi]\!]^c) \subseteq [\![\psi]\!]^c$ (the best ϕ-worlds are ψ-worlds), that $[\![\phi]\!]^{c,i} = 1$, and that $\neg\phi$ implies $\neg\psi$. Then, by Definition 2, $[\![(if\ \phi)(\Box \psi)]\!]^{c,i} = 1$. But since $\neg\phi$ implies $\neg\psi$ (i.e., $[\![\neg\phi]\!]^c \subseteq [\![\neg\psi]\!]^c$) and $D_c(F_c(i)) \subseteq [\![\neg\phi]\!]^c$, $D_c(F_c(i)) \subseteq [\![\neg\psi]\!]^c$. So $[\![\Box \psi]\!]^{c,i} = 0$. Indeed, $[\![\Box\neg\psi]\!]^{c,i} = 1$. □

---

[6] See also Lewis (1975), Kratzer (1991). For indicatives whose consequents lack overt modals (e.g., *if he has his umbrella, it's raining*), Kratzer posits a covert quantifier at LF. For critical discussion of this part of Kratzer's proposal, see Gillies (2010, Sect. 9).

This countermodel appealed to in this proof is, notice, similar to the one used in proving invalidity of modus ponens for the wide-scoper.

But, promisingly, there is no immediate way to recreate the problem that afflicted wide-scoping for the restrictor semantics, namely predicting the sentences in (3) inconsistent, on the assumption that $F_c \subseteq [\![in\_A \vee in\_B]\!]^c$. The wide-scoper's problem stems from three properties of the *unrestricted* domain $D_c(F_c)$, namely:

- $D_c(F_c(i)) \subseteq [\![in\_A \supset block\_A]\!]^c$
- $D_c(F_c(i)) \subseteq [\![in\_B \supset block\_B]\!]^c$
- $D_c(F_c(i)) \subseteq [\![in\_A \vee in\_B]\!]^c$

In a restrictor semantics, however, different restrictors for $\square$ induce *three distinct domains* of quantification, so that, in lieu of three conditions on $D_c(F_c)$, the same suppositions yield the following:

- $D_c(F_c(i) \cap [\![in\_A]\!]^c) \subseteq [\![block\_A]\!]^c$
- $D_c(F_c(i) \cap [\![in\_B]\!]^c) \subseteq [\![block\_B]\!]^c$
- $D_c(F_c(i)) \subseteq [\![in\_A \vee in\_B]\!]^c$

All good, then? No. For the wide-scoper's problem reemerges if we adopt a very natural constraint on the selection function. (Here and throughout, W refers to the universe of possibilities.)

> STABILITY
> $\forall p, p' \subseteq W$: if $p' \subseteq p, i \in p'$, and $i \in D_c(p)$, then $i \in D_c(p')$

Stability has it that if a possibility $i \in p$ is best, i remains best in any strengthening (contraction) of p. Stability is extremely *prima facie* plausible. If a possibility has enough (with respect to other possibilities in a set p) good-making features, then it does not cease having enough good-making features with respect to a contraction of p. Contracting p, if anything, *reduces the possibility's competition*. Denying Stability is, at first glance, rather like denying that the best restaurant in Manhattan (which happens to be located in SoHo) must also be the best restaurant in SoHo. (We will return to this point in Sect. 5.)

Stability, together with the basic Restrictor Semantics, predicts the sentences in (3) inconsistent, on the assumption that $F_c(i) \subseteq [\![in\_A \vee in\_B]\!]^c$.

Proof Suppose otherwise. Choose $i' \in D_c(F_c(i))$. By Realism, $i' \in F_c(i)$.

> Since $F_c(i) \subseteq [\![in\_A \vee in\_B]\!]^c$, $i' \in F_c(i) \cap [\![in\_A]\!]^c$ or $i' \in F_c(i) \cap [\![in\_B]\!]^c$.
> So, by Stability, $i' \in D_c(F_c(i) \cap [\![in\_A]\!]^c)$ or $i' \in D_c(F_c(i) \cap [\![in\_B]\!]^c)$.
> By Definition 2, $D_c(F_c(i) \cap [\![in\_A]\!]^c) \subseteq [\![block\_A]\!]^c$ and $D_c(F_c(i) \cap [\![in\_B]\!]^c) \subseteq [\![block\_B]\!]^c$.
> So $D_c(F_c(i)) \subseteq [\![block\_A]\!]^c \cup [\![block\_B]\!]^c$.
> So, by Definition 2, $[\![\square(block\_A \vee block\_B)]\!]^{c,i} = 1$. Contradiction. $\square$

That is a disaster.[7] What has gone wrong? These are the possibilities:

---

[7] I don't mean to suggest K&M are unaware of this. Quite the opposite—they explicitly reject Stability (K&M, p. 133). My problems with their account are twofold. (1) They give undue prominence to the denial

– The Restrictor Semantics
– The constraints on D_c (Realism, Stability)

For our purposes, Realism is nonnegotiable: deontic modals (of the sort that bear on deliberation about what to do) are quantifiers over relevant possibilities. Insofar as Stability is plausible, we should be on the lookout for a different semantics for modals and indicatives (indeed, we'll consider one in the next section). The take-home point, for now, is that endorsing a genuine denial of modus ponens is, by itself, insufficient for avoiding the puzzle.

3.2 Shifty semantics

Invalidating modus ponens is necessary for resolving the puzzle. Having rejected wide-scoping and (tentatively) the Restrictor analysis, we have committed to narrow-scoping the modals of (1b) and (1c), so that they remain *in situ* at LF.

That is a hard row to hoe, since the major non-Restrictor analysis of indicatives— the "Variably Strict" analysis (Stalnaker 1968; Lewis 1973)—validates modus ponens. The Stalnaker-Lewis analysis interprets indicatives relative to a similarity ordering $\leq_i$ that is weakly centered on i:

WEAK CENTERING
$\forall p \subseteq W : i$ is $\leq_i$ minimal in p (i.e., $\neg \exists j \in p : j \leq_i i \wedge i \not\leq_i j$)

Weak centering codifies the intuition that, if i is a ϕ-world, there are no ϕ-worlds closer to i than i itself. Let $\rightsquigarrow$ represent a variably strict conditional operator at LF. On the Stalnaker-Lewis analysis, an indicative of the form $\ulcorner \phi \rightsquigarrow \psi \urcorner$ is true iff all the closest ϕ-worlds are ψ worlds[8]:

Definition 3 $[\![\phi \rightsquigarrow \psi]\!]^{c,i} = 1$ iff $\forall j \in [\![\phi]\!]^c : j$ is $\leq_i$-minimal in $[\![\phi]\!]^c \Rightarrow [\![\psi]\!]^{c,j} = 1$.

Weak Centering ensures variably strict conditionals will validate modus ponens.

Proof Suppose $[\![\phi \rightsquigarrow \psi]\!]^{c,i} = [\![\phi]\!]^{c,i} = 1$. Then, by Definition 3, for any $j \in [\![\phi]\!]^c$ that is $\leq_i$-minimal in $[\![\phi]\!]^c$, $[\![\psi]\!]^{c,j} = 1$. By Weak Centering, i is $\leq_i$-minimal in $[\![\phi]\!]^c$. By supposition, $[\![\phi]\!]^{c,i} = 1$. So $[\![\psi]\!]^{c,i} = 1$. □

So the correct semantics for indicatives *cannot* be the Stalnaker-Lewis semantics.

A promising alternative—the one, in fact, endorsed by K&M—is based on the aforementioned Ramsey test.[9] The Ramsey test for the truth of (*if* ϕ)($\Box \psi$) involves checking whether $\Box \psi$ holds *throughout* a relevant body of information carrying the information that ϕ. Some preliminary definitions:

Footnote 7 continued

of modus ponens in their explanation of the puzzle's resolution; the rejection of Stability is very much in the background (indeed, they do not actually offer any proofs or arguments that it leads to paradoxical consequences). (2) They do not offer a persuasive philosophical motivation for their denial of Stability, so their account is explanatorily incomplete. I elaborate on these points in Sect. 4.

[8] Once again, we make the Limit Assumption. This does not affect the proof.

[9] For similar proposals, also based on the Ramsey test, see Yalcin (2007), Gillies (2010).

Definition 4 ϕ is *true throughout a domain p* at c (notation: $p \models_c \psi$) iff $p \subseteq [\![\phi]\!]^c$

Definition 5 A *shifted context* $c + \phi$ is just like c, except $\forall i : F_{c+\phi}(i) := F_c(i) \cap [\![\phi]\!]^c$

With these definitions in hand, the Ramsey test seems to yield a simple semantics for indicatives. Informally, $(if\ \phi)(\Box \psi)$ is true at c just in case $\Box \phi$ is true throughout $F_{c+\phi}$ (i.e., true throughout the original domain, incremented with the information that ϕ). Formally, letting $\hookrightarrow$ represent a shifty conditional operator at LF:

Definition 6 $[\![\phi \hookrightarrow \psi]\!]^{c,i} = 1$ iff $F_{c+\phi}(i) \models_{c+\phi} \psi$

This first pass, as K&M note, is not quite right. As defined, $F_{c+\phi}$ does not necessarily bear the information that ϕ. There are, in fact, cases where $F_{c+\phi}(i) \models_c \neg\phi$, i.e., where $F_{c+\phi}$ bears the information that $\neg\phi$ (K&M, p. 135)! Consider an epistemic operator *might*, which says that its sentential complement ("prejacent") is compatible with the relevant information, and the conditional in (4):

(4) $(if\ in\_A \wedge might(\neg in\_A))(\Box \psi)$
   ($\approx$ If they're in A but our information doesn't entail that, we should realize ψ)

Let χ abbreviate (4)'s antecedent. Since $F_{c+\chi}(i) \subseteq [\![in\_A]\!]^c$, $might(\neg in\_A)$ is *false*, rather than true, throughout $F_{c+\chi}(i)$. So, $F_{c+\chi}(i) \models_{c+\chi} \neg\chi$. It offends common sense (as embodied in the Ramsey test) to use such a domain in evaluating (4).

The fix is simple: *ensure the domain bears the information that the antecedent is true*. This is done with a tweak to Definition 5 (K&M, p. 136; cf. Yalcin 2007, p. 998).

Definition 7 A *shifted context* $c + \phi$ is just like c, except $\forall i : F_{c+\phi}(i)$ is the largest subset of $F_c(i)$ such that $F_{c+\phi}(i) \models_{c+\phi} \phi$

This, with the semantics in Definition 6 (hereafter, the "Shifty Semantics"), is more consonant with the intuitions behind the Ramsey test: evaluating $(if\ \phi)(\Box \psi)$ at a context c involves checking whether $\Box \psi$ holds throughout a body of information that bears the information that ϕ (and is as similar to the original information as possible).[10]

And what of modus ponens? As with the Restrictor treatment, it comes out *invalid* on the Shifty semantics. The reason, informally, is that the truth of a shifty conditional $\phi \hookrightarrow \psi$ at $\langle c, i \rangle$ implies only that ψ holds throughout a domain carrying the information that ϕ. But the "basic" domain, F_c, may not bear the information that ϕ. Although this makes no difference if ψ's truth is context-independent (in which case, the fact that $[\![\psi]\!]^{c+\phi,i} = 1$ will imply that $[\![\psi]\!]^{c,i} = 1$), it can make a *major* difference when ψ's truth is dependent on contextually relevant information (if, for instance, ψ is of the form $\ulcorner \Box \chi \urcorner$).

In short, the Shifty Semantics for indicatives validates an ersatz, but not a genuine, form of modus ponens.

---

[10] Definition 7 presupposes (wrongly) there is always just one such subset (K&M, p. 136). Although this leads them to adopt a more complicated notion of a shifted context than embodied in Definition 7, nothing here turns on this, and so we will prefer the simpler definition.

GENUINE MODUS PONENS

$[\![(if\ \phi)(\psi)]\!]^{c,i} = [\![\phi]\!]^{c,i} = 1$ implies $[\![\psi]\!]^{c,i} = 1$

ERSATZ MODUS PONENS

$[\![(if\ \phi)(\psi)]\!]^{c,i} = [\![\phi]\!]^{c+\phi,i} = 1$ implies $[\![\psi]\!]^{c+\phi,i} = 1$

Proof The invalidity of genuine modus ponens on the Shifty Semantics is obvious enough. So we prove the validity of erstaz modus ponens on the Shifty Semantics. Suppose that $[\![(if\ \phi)(\psi)]\!]^{c,i} = [\![\phi]\!]^{c+\phi,i} = 1$. By Definition 6, $F_{c+\phi}(i) \models_{c+\phi} \psi$. So, by Definition 4, $F_{c+\phi}(i) \subseteq [\![\psi]\!]^{c+\phi}$. By the Reflexivity of Relevance, $i \in F_c(i)$. Since $i \in [\![\phi]\!]^{c+\phi}$ and $i \in F_c(i)$, $i \in F_{c+\phi}(i)$, by Definition 7. And since $F_{c+\phi}(i) \subseteq [\![\psi]\!]^{c+\phi}$, $i \in [\![\psi]\!]^{c+\phi}$. □

This feels like progress! The Ramsey readings of (1b) and (1c), to which we appealed to motivate the puzzle in Sect. 2, do seem to license only an ersatz form of modus ponens. Drawing a conclusion about what shaft to block isn't licensed at the basic context. But adding some information to the context changes things: once our information settles the miners' location, it settles which shaft to block. The Shifty Semantics for indicatives seems tailor-made to capture such intuitions.

So, all good, then? Sadly, no. We are saddled with *the same variant of the puzzle as afflicted the Restrictor Semantics*, if we endorse an adapted version of the Stability constraint on the selection function.

MONOTONICITY

If $F_{c'}(i) \subseteq F_c(i)$, c and c' are otherwise identical, and $j \in F_{c'}(i)$, then $j \in D_c(F_c(i))$ implies $j \in D_{c'}(F_{c'}(i))$

Monotonicity is a notational variant of Stability. Stability says that, if a possibility is best in a larger domain, it remains best when we reduce its competition (contract the domain); how possibilities are ranked by the deontic selection function is *independent of contractions of the domain*. This is just what Monotonicity says explicitly.

Monotonicity, together with the Shifty Semantics (and the usual quantificational semantics for deontic modals in Definition 1) predicts the sentences in (3) inconsistent, on the assumption that $F_c(i) \subseteq [\![in\_A \vee in\_B]\!]^c$.

Proof Suppose (3) is true at $\langle c, i \rangle$. Choose $i' \in D_c(F_c(i))$. By Realism, $i' \in F_c(i)$.

By Definition 7, since $F_c(i) \subseteq [\![in\_A \vee in\_B]\!]^c$, $i' \in F_{c+in\_A}(i)$ or $i' \in F_{c+in\_B}(i)$.
So, by Monotonicity, $i' \in D_{c+in\_A}(F_{c+in\_A}(i))$ or $i' \in D_{c+in\_B}(F_{c+in\_B}(i))$.
By the Reflexivity of Relevance, $i \in F_c(i)$.
Since $i \in F_c(i)$ and $F_c(i) \subseteq [\![in\_A \vee in\_B]\!]^c$, $i \in [\![in\_A]\!]^c$ or $i \in [\![in\_B]\!]^c$.
So, since *in_A / in_B* are context-invariant, $i \in [\![in\_A]\!]^{c+in\_A}$ or $i \in [\![in\_B]\!]^{c+in\_B}$.
So, using Ersatz Modus Ponens, $i \in [\![\Box block\_A]\!]^{c+in\_A}$ or $i \in [\![\Box block\_B]\!]^{c+in\_B}$.
So, by Definition 1, $D_{c+in\_A}(F_{c+in\_A}(i)) \subseteq [\![block\_A]\!]^{c+in\_A}$, or $D_{c+in\_B}(F_{c+in\_B}(i)) \subseteq [\![block\_B]\!]^{c+in\_B}$.
So, either $i' \in [\![block\_A]\!]^{c+in\_A}$ or $i' \in [\![block\_B]\!]^{c+in\_B}$.
So, since *block_A / block_B* are context-invariant, $i' \in [\![block\_A]\!]^c$ or $i' \in [\![block\_B]\!]^c$.

So $D_c(F_c(i)) \subseteq [\![block\_A]\!]^c \cup [\![block\_B]\!]^c$.
So $[\![\Box(block\_A \vee block\_B)]\!]^{c,i} = 1$. Contradiction. $\qquad\qquad\square$

The proof is involved, but the driving idea is simple (indeed, the same as the idea behind the proof of the same result for the Restrictor Semantics): if i' is one of the best worlds at c, i' must (given Monotonicity) remain best at c-plus-some-information-about-the-miners'-location. But, at an informationally enriched context, we have it that the best worlds all involve blocking the shaft we know the miners are in. So, at i' (and the best worlds at the basic context, generally), we block at least one of the shafts. Despite a promising start, the basic structure of the problem is unchanged.

4 What an explanation requires

We've argued, with K&M, that solving the miner puzzle requires giving an explanation of the consistency of the sentences in (1) and the consistency of those in (3) when the domain entails that the miners are in A or B. Any such explanation demands rejecting a genuine modus ponens rule for genuine indicatives. That is striking.[11]

But it is not enough. The lesson of the prior section is that this cannot be the sole element in an explanation of the consistency of (1) and the consistency of (3). For the very same puzzle arises on analyses of indicative conditionals that invalidate modus ponens. Insofar as these puzzles are the same, we should, other things being equal, be on the lookout for a *comprehensive solution*. Since the problem arises even on accounts that presuppose the invalidity of modus ponens, rejecting the validity of modus ponens cannot be the comprehensive solution we are seeking. This section describes the features such a solution should have, and argues that K&M's account lacks them.[12]

Desideratum 1 A *Theory-External* Rationale. K&M, as I noted earlier, recognize that rejecting modus ponens does not go far enough. They note (although they do not give arguments) that their semantics for indicatives and deontic modals does not handle the explananda unless they reject Stability, but they do not give this very much attention. They *do*, of course, reject Stability (see K&M, p. 133). There is a real worry, however, that this rejection is unmotivated and *ad hoc*—that the proper reaction to the persistence of the puzzle is to opt for a new treatment of indicatives and modals, rather than to reject a potentially independently well-motivated constraint on the deontic selection function for the sake of making the Shifty Semantics consonant with the data. Although K&M do supply a brief rationale for the rejection of Stability, that rationale is (as I'll argue in the next section) somewhat unconvincing (especially in light of the various considerations that seem to recommend a Stability constraint on selection functions). That should be cause for worry.

---

[11] Thony Gillies tartly notes: "You have to troll some pretty dark corners of logical space for deniers of modus ponens" (Gillies 2010, p. 14).

[12] For similar, independently developed critiques of K&M's account, see Silk (ms) and Cariani et al. (ms). Both of these references defend resolutions of the Miner Puzzle that are, I believe, similar in inspiration to that defended here. Although similar inspiration, there are fundamental differences in implementation between these approaches and mine (and, I believe, strong reasons for preferring the approach taken here). I regret that I lack the space to explore those differences here.

Desideratum 2 Explain Consistency by Explaining Truth. My intuition, like K&M's, of the consistency of the sentences in (1) and (3) stems from an intuition of their *truth* at the described context. So it would be desirable to have an account which explained the consistency of (1) and (3) by doing justice to these intuitions—by giving a semantics which predicted the truth of (1) and (3) at the described context. Rejecting Stability, even if well-motivated, generates a mere prediction of consistency. Insofar the intuition of consistency is *explained by* the intuition of truth, a more ambitious account than the one on offer from K&M would be explanatorily preferable. (If the semantics supplied by that sort of explanation came with a plausible, theory-external rationale, that would be gravy. That is just what our account will ultimately provide.)

A natural response on K&M's behalf is that it is not the job of the *semantics*, as such, to predict the truth of substantive deontic claims like those in (1) and (3).[13] So stated, this is, I think, is quite right. Such facts, if they are facts at all, are facts about obligation (or normativity, or perhaps rationality), rather than semantic facts. In reply, however, I submit that it is a general condition of adequacy on a semantics for indicative conditionals and deontic modals that it explain how deontic (or normative, or rational) concerns—facts, roughly, about what is deontically (or normatively, or rationally) preferred—interact with contexts (specifically, context-dependent facts about what information is relevant) to yield truth-values, at realistic contexts, for substantive claims about what to do (conditional and otherwise). In the miner case, specifically, a semantics should explain how these variables interact to yield the *expected* truth-values, at the described contexts, for the claims in (1) and (3). In that case, I propose, this means stating a semantics that does justice to each of the following near-platitudes:

– The preferred outcome is one in which all the miners are saved.
– Our ignorance of a relevant fact (namely, the miners' location) makes it the case that, although the preferred outcome is one in which all the miners are saved, it is not the case that we ought to realize this outcome.
– It can nevertheless be the case that we ought to realize this outcome (by blocking shaft X) *on the supposition* that the miners are in shaft X.

For a semantics to 'do justice to' these claims, I submit that it should explain:

– How the computation of truth-values of the *ought*s in (1) and (3)—in particular (1d) and (3c)—is sensitive to facts about *preferred outcomes*.
– How context-dependent *limitations on information* can properly impact an agent's pursuit of preferred outcomes (hence, how such limitations can affect the interpretation of linguistic constructions expressing how an agent ought properly to pursue preferred outcomes—namely, descriptions of that agent's obligations).
– How *making a supposition* that ϕ can also properly impact an agent's deliberations about what to do in ϕ circumstances (hence, how such supposing that ϕ can affect the interpretation of linguistic constructions expressing obligation conditional on ϕ), without correspondingly impacting her deliberations about what to do when that supposition is discharged.

---

[13] I'd like to thank an anonymous reviewer for pressing this point.

While K&M do make some suggestive gestures at such explanations (as we shall see in Sect. 5), these gestures do not amount to any sort of account (nor do they attempt to develop the sort of semantic apparatus within which such an account might be stated). That is an explanatory gap in their account. To emphasize: the major "criticism" of K&M's account represented by this point (as well as the points advanced in our discussions of Desiderata 1 and 3) is one of *explanatory incompleteness*. K&M could, I expect, accept a story much like the one on offer in this paper without committing themselves to rejecting any central part of their account (save for the meta-theoretical claim that their account constitutes an explanation of the Miner Puzzle).

It bears emphasizing that neither the explanatory gap I am highlighting nor its resolution are at all trivial. As we've seen and will continue to see, it, rather than the failure of modus ponens for indicatives, is at the *heart* of the Miner Puzzle (whereas the rejection of modus ponens for indicative conditionals is orthogonal). Its resolution will involve (i) developing a view about the interplay between considerations of value and relevant information in deliberation about what to do (including deliberation about what to do in so-and-so circumstances), (ii) implementing this view in a semantics for constructions expressing claims about what to do (including constructions expressing claims about what to do in so-and-so circumstances). (I return to these points in Sect. 5.)

The Miner Puzzle, in short, is deeply puzzling. K&M's discussion leaves central portions of its explanation unexplored. So an account which remedies its explanatory deficiencies is to be preferred. Such an account is just what this paper will ultimately attempt to develop.

Desideratum 3 Explain the Weak / Strong Asymmetry. There is a firm intuition that the sentences in (1) and (3) are all consistent. However, there is also a firm intuition—fairly reliably corroborated by informants—that the sentences in (5) are *inconsistent* in the case as described.[14]

(5)   a.   They're in A or B.
      b.   If they're in A, we must [have to] block A.
      c.   If they're in B, we must [have to] block B.
      d.   We may block neither shaft.

Compare, also, the following dialogues.

(6)   a.   If they're in A, we should block A. And if they're in B, we should block B.
      b.   But what do you think we should do?
      a.   Block neither.

(7)   a.   If they're in A, we must block A. And if they're in B, we must block B.
      b.   But what do you think we must do?
      a.   #/?Block neither.

---

[14] The most common reaction to (5) is a sort of feeling of oddness. I take that to be prima facie evidence of a judgment of inconsistency. I do not want to suggest that the intuitions are clear, or that the claimed oddness has the status of a datum. Some, in fact, do not hear the oddness (although the large majority of my informants do). The felt oddness is sensitive, as these things tend to be, to shifts in focus/emphasis and the manner in which the context is described. While I do think it is robust, and that it supports a judgment of inconsistency, defending these claims fully is beyond the scope of this essay.

Why do these asymmetries matter? To this point, we've considered only examples invoking the *weak* deontic necessity modal *should* (although the same points would have held for *ought*). But, as these examples illustrate, shifting to examples invoking *strong* deontic necessity modals (*must*, *have to*, …) changes things.[15] That would be surprising, if our solution to the original puzzle traded only on the denial of modus ponens. For (3) and (5) differ in the strength of their modals, but are otherwise identical; if (3a) contains, at LF, a shifty conditional operator that fails to validate modus ponens, it would stand to reason that the same such operator occurs in (5b).

That a solution requires more than just denying modus ponens should, by now, be a familiar point. But there is, nevertheless, something striking and important about the fact that strengthening the modals affects consistency intuitions in this way. It suggests:

– While indicatives with weakly modal consequents do not generally allow detachment of their consequents, on the supposition that their antecedents are true, indicatives with strongly modal consequents *do* allow *some* sort of detachment. (We will remain noncommittal about what sort, exactly.)
– Whatever conditions on the semantics of indicatives with weakly modal consequents we add to the denial of modus ponens to explain the consistency of (1) and (3), *cannot hold for indicatives with strongly modal consequents*.

Weak and strong necessity modals must, in short, display different interactions with relevant information. A complete solution to the puzzle should not, then, just succeed in predicting the consistency of (1) and (3). It should predict the consistency of (1) and (3) and the inconsistency of (5) as a function of the distinct ways in which weak and strong necessity modals, respectively, interact with relevant information.

Here, then, is another central dimension of the puzzle that K&M's account fails to explain. We will do our best to fill in this lacuna (as well as those associated with the other desiderata) in what follows. There is, it turns out, an illuminating connection between the strength of a deontic modal and the way it interacts with relevant information. Indeed, modal strength (more specifically, the nature of the different kinds of deontic modality invoked by different kinds of deontic necessity modals) actually *explains* differential sensitivity to relevant information.

5 Stability, information-dependence, and practical reasons

In this section, we will take a closer look at Stability (and thereby at Monotonicity) to see whether there is a good case against adopting it as a constraint on the deontic selection function.

STABILITY
$\forall p, p' \subseteq W$: if $p' \subseteq p, i \in p'$, and $i \in D_c(p)$, then $i \in D_c(p')$

---

[15] Strong and weak necessity modals are so-called because the former are thought to asymmetrically entail the latter. For a sophisticated recent discussion of weak necessity, see von Fintel and Iatridou (2008).

To preview, we look at K&M's brief rationale for rejecting Stability, and argue that it fails. We consider some ways that the rationale might be modified, and make some tentative gestures at a semantics for strong and weak deontic necessity modals which implements these insights.

Stability, I suggest, is a reasonable constraint on the "basic" deontic selection function. But, drawing on recent work on weak deontic modality by von Fintel and Iatridou (2008), we argue that the domains over which weak deontic necessity modals quantify are determined, *not just* by the basic deontic selection function, but by other selection functions which rank worlds according to "secondary" desiderata (whereas the domains for strong deontic necessity modals are determined just by the basic deontic selection function). We retain Stability (and the Shifty Semantics for indicatives), but are able to avoid the relevant difficulties *by revising the truth-conditions for weak deontic necessity modals*.

5.1 Locating information-dependence

In this section, I want to advance a tentative argument for Stability with respect to the "basic" deontic selection function (more on basic-ness below). As a preliminary to this, it is useful to get clear on just what sort of constraint Stability amounts to. That is somewhat easier said than done: Stability is subject to various plausible glosses, which makes evaluating its plausibility somewhat fraught. Taking for granted that D_c selects the deontically *best* (or, more strongly, *ideal*) possibilities from a set, both of these glosses are live options.

– *Gloss One.* If a possibility i is ideal in a set of possibilities p (i.e., i has enough, relative to the other possibilities in p, good-making features), it does not cease having enough good-making features in a contraction of p.
– *Gloss Two.* What makes a possibility ideal in a set of possibilities p does not depend on the information contained in p. That is to say: good-making features are invariant under information-acquisition.

Gloss One, I've suggested, casts Stability in an extremely plausible light. Indeed, a parallel constraint on *choice functions* (known variously by the names 'Independence of Irrelevant Alternatives' (IIA), 'Principle α', and 'Property α') figures prominently in the philosophical and economic literature on preference and rational choice. Choice functions are basically analogous to our deontic selection function: both map a preference-ordering R and range of alternatives X ordered by R to the set of $x \in X$ such that no $y \in X$ is strictly R-better than x (i.e., $\{x \in X : \neg\exists y \in X : yRx \land \neg xRy\}$). What differentiates choice functions and selection functions, most saliently, is the fact that, once a preference ordering is fixed, the former are defined on sets of outcomes (states of affairs), rather than sets of maximally specific possibilities. Amartya Sen calls IIA a "very basic requirement of rational choice… [It] states that if the world champion in some game is a Pakistani, then he must also be the champion in Pakistan" (Sen 1969, p. 384).[16] If this is a solid rationale for IIA for choice functions (and it is), it is

---

[16] There are purported counterexamples to IIA, none of which are very persuasive or relevant to our discussion here. For discussion, see Hansson (2005).

a prima facie solid rationale for Stability for deontic selection functions. At the very least, the denier of Stability (who accepts IIA for choice functions) incurs a real burden to explain the basis on which choice functions and deontic selection functions should be differentiated in this respect.

Gloss Two might seem to provide such a basis. For, against this gloss of Stability, we might think that ends, values, concerns, etc.—the criteria by which the ordering on possibilities is determined—*properly depend on relevant information* (in our case, the information carried in the set of alternatives on which the selection function is defined). And, in fact, this is exactly what K&M would like to suggest. Here is K&M's case against the Stability constraint on deontic selection functions.

> [I]t is not just the set of ideal worlds that varies as the information state is shifted, but also the *ranking* of worlds as more or less ideal. A world may be more ideal than another relative to one information state and less ideal than it relative to another. For example, a world in which both shafts are left open may be more ideal than one in which shaft A is closed relative to a less informed state, but less ideal relative to a more informed state. (p. 133)

The most straightforward reading of this claim seems to me *plainly false* (although, as we will see just below, there is reason to think K&M do not actually intend that reading.[17]) Worlds where all of the miners are saved are best (ideal); worlds where only nine are saved are strictly less ideal. While choosing the path of no risk (i.e., doing nothing) is, I think, the correct decision in this case, it is *not* because worlds in which nine miners are saved are ideal, even from a standpoint of subjective uncertainty. It is because our epistemic state fails to put us in a *position to implement*, with sufficient certainty, those ends, values, concerns, etc., which determine which possibilities are ideal. Indeed, this is just what it means to say that our predicament is *risky*. If there is unexpected variation in the quantificational domains of weak deontic necessity modals here, it would not seem to come from ideality varying with information.

This is *not* to deny, generally, that ends that are constitutive of the ideal can depend on our information. It is a familiar, if perhaps controversial, point—one dating, indeed, to Aristotle's discussions of the practical syllogism (cf. Wiggins 1975; Burnyeat 1980)— that not all practical reasoning involves using information to determine how to implement given ends. Some practical reasoning is plausibly *rule-case*, rather than *means-end*. Roughly, means-end practical reasoning involves thinking about what actions are *causally conducive* to a given end (thinking about what is instrumentally valuable), while rule-case reasoning involves thinking about what sorts of things are constitutive of something that is an end-in-itself (thinking about what is intrinsically valuable). The conclusions of rule-case reasoning, insofar as they are conclusions about what is intrinsically rather than instrumentally valuable, bear directly on what is ideal. And rule-case reasoning can be affected by information (about what is constitutive of, rather than conducive to, the ideal). If, for instance, eudaimonia is constitutive of the

---

[17] To preview, it is possible that K&M are using a different notion of ideality: one (perhaps) on which i is ideal iff, for any j, one does as well relative to *what one should do* in i as in j, where what one should do depends on both first-order ends (saving life) and second-order ends (pursuing only pursuit-worthy first-order ends). I will address this possibility in greater detail shortly.

ideal, and we learn that virtue is constitutive of eudaimonia, then virtue might thereby become one of the ends that is constitutive of the ideal. In no way do I want to be interpreted here as ruling out the sort of information-dependence associated with rule-case reasoning.

Notice, though, that this is not the sort of information-dependence that the miner case involves. Learning the miners' location puts us in a position to take action that is causally conducive to the ideal (saving everyone). But it does not bear in any way on what *is* ideal—on what the best course of action, given a set of ends, is. Instead, it seems to bear, very roughly, on *whether some ideality-contributing consideration is allowed to influence our practical reasoning in a certain way*. Ideality-contributing considerations are allowed to influence our practical reasoning in the relevant way only when the body of information puts us in a position to take (or advise taking) an action that is causally conducive to the realization of that consideration.

K&M might reply that this does not necessarily impugn their account. Aside from (what we might call) *first-order* or *substantive ends*—ends whose realization by a state of affairs contributes to its intrinsic value—we should recognize *second-order* or *procedural ends*—ends which specify, roughly, how best to go about selecting an action for the sake of realizing substantive ends. What is best, or most choice-worthy, is a function of *both* first- and second-order ends.[18] We routinely make such distinctions in

---

[18] This possibility might seem to be supported by their brief discussion of the nature of deontic ideality. They note, for instance, "It is natural to think that the species of deontic ideality relevant to our efforts to save the miners depends somehow on choice" (p. 132). In their Footnote 28, they go on to explore one way of cashing out this dependence, on which a deontic sentence $\Box\phi$ requires for its truth that there be some available action α such that the relevant information entails that, if α is done, ϕ. In a slogan, it cannot be that we ought to realize ϕ if there is no available action whose performance, it's known, realizes ϕ (or, as they put it, "*ought* implies *can choose*"). And this would seem to support attributing to them a view on which what is deontically ideal is (somehow) a function both of what outcomes are desirable (first-order ends) and epistemic constraints on the pursuit of desirable outcomes (second-order ends).

Notice, however, that such a principle *has no bearing on deontic ideality* per se. Indeed, one can accept it as a property of a deontic claim $\Box\phi$ that its truth requires that ϕ is choosable, without thinking this property follows from (or, indeed, is in any way related to) properties of the deontic selection function (e.g., Stability). Here is a helpful analogy: one might plausibly think that actions that are decision-theoretically rational must be choosable, in roughly K&M's sense, *without* thinking that this requirement is explained by the properties of *expected utility functions*. Indeed, decision theorists typically build choosability into their *descriptions of decision matrices*, rather than into their descriptions of expected utility functions; expected utility functions are *defined* only for actions that figure in the description of the decision matrix (where one requirement on figuring in such descriptions is, typically, something not far removed from choosability). Analogously, a semanticist might elect to encode choosability as a lexicalized *presupposition* of deontic necessity claims. Alternatively, choosability might be regarded as a semantically independent condition on the truth of deontic necessity claims, so that $\Box\phi$ is true just if the best worlds are ϕ-worlds and ϕ is chooseable relative to the relevant information (as, e.g., *belief* and *safety* are semantically independent conditions on the truth of knowledge ascriptions). Significantly, neither course will give the theorist any direction in describing the properties of the deontic selection function.

Even if K&M endorse the notion that choosability must follow from properties of the deontic selection function, the principle per se provides no clear guidance in theorizing about what the deontic selection function must *be like* in order for choosability to be validated (or how the relevant properties of the deontic selection function follow from the interaction of first- and second-order ends). (Similarly, because of the independence of choosability requirements on decision-theoretic verdicts and expected utility functions, endorsing a choosability requirement on decision-theoretic verdicts gives no guidance in formulating a correct account of the properties of expected utility functions.) It is, of course, open to K&M to embrace the general sort of picture of deontic ideality I am developing on their behalf here. Because of the independence

other domains. In, e.g., normative decision theory, a basic distinction is made between substantive ends (which are supplied by an agent's utility function) and procedural ends (which are supplied by the decision theorist, and specify what constraints a choice function relating a utility and probability function to chosen actions must satisfy *in order to count as practically rational*). There is a clear sense in which the actions recommended by a rational choice function under conditions of subjective uncertainty are best or most choice-worthy, even though they may not, given the actual state of the world, maximize utility. They promote the end of maximizing *expected* utility, and it is something of a decision-theoretic mantra that an action is best or most choice-worthy if and only if it maximizes expected utility.

All that is right. But for it to help K&M, it must supply some sort of reason for building epistemic constraints *into* the deontic selection function. But it does not do this (just as similar considerations supply no reason to represent the end of maximizing expected utility in the utility function itself). Indeed, taking the decision-theoretic analogy seriously, it seems to recommend the introduction of some sort of *independent theoretical apparatus* taking some specification of (i) substantive, first-order ends and (ii) an information-state as its input, and generating a modified ordering on possibilities as output.[19] We could, of course, conflate what is "best" or "ideal" with what is minimal *with respect to the modified ordering*. Semantically, this would amount to restyling the deontic selection as selecting, not worlds that are best *sans phrase*, but with respect to some mixture of first- and second-order ends—to replacing a selection function determined by first-order ends with one determined by first- and second-order ends, and an information state. Such a selection function would, in fact, not respect Stability.

This, however, is problematic. If *the* deontic selection function fails to respect Stability, then we will predict any modal that makes use of the deontic selection function (strong deontic necessity modals included) to exhibit the phenomena associated with Stability violations. So there will be trouble explaining the weak/strong asymmetry (Sect. 4). *Even if* this can be avoided, a theoretical point still stands: the theorist should take care to avoid running together theoretical notions (and pieces of theoretical apparatus) that, like actual and expected utility, a theory ought not to run together. Distinguishing these pieces of apparatus fills explanatory gaps that would otherwise be left open. An explanatorily adequate account of the puzzle will furnish (i) an account of the relationship between the good and the choice-worthy, (ii) a semantics for weak necessity modals which exploits this account. The rest of this paper aims to provide both (i) and (ii).

Footnote 18 continued

of choosability and deontic ideality, however, their paper cannot strictly be regarded as putting forth any sort of account of the semantic contribution of first- and second-order ends in the truth of deontic claims (nor does it give any guidance, per se, about how to go about constructing such an account). We will defend such an account for deontic selection functions in Sect. 6.

[19] There are, additionally, linguistic reasons (having to do with counterfactual morphology of weakly modal constructions in, e.g., French) for thinking weak deontic necessity modals must make use of two orderings, one corresponding to primary (first-order) ends, the other to secondary (second-order) ends. See esp. von Fintel and Iatridou (2008) and the discussion in Sect. 6.

5.2 Characterizing information-dependence

In this section, I will *very roughly* sketch one picture of how relevant information properly bears on practical reasoning—one inspired by the account of epistemic constraints on treating a consideration as a reason of Hawthorne and Stanley (2008). It is good to have in view a reasonably well-motivated picture of how information affects the appropriateness of allowing an ideality-contributing consideration to influence practical reasoning before stating a formal representation of that relationship as part of our semantics. That is what we're after here. Caveat lector: we won't offer any arguments, per se, in favor of this exact picture. We are simply after something sufficiently plausible and concrete to be of theoretical use.

Let's begin by distinguishing, roughly, between *basic ends* and *practical ends*. Basic ends are first-order ends—the ends that provide overarching guidance to practical deliberation (and which partition possibilities into ideal and non-ideal). Practical ends are *reason-giving ends*, in the sense that, if p is a practical end and α and β are available actions:

– If α realizes p, it is appropriate to treat the proposition that α realizes p as a reason for performing α.
– if β fails to realize p, it is appropriate to treat the proposition that β fails to realize p as a reason against performing β.

Clearly, there is some sort of connection between practical ends and basic ends—some sort of connection between first-order concerns and those considerations it is appropriate to treat as reasons for or against various courses of action. Indeed, we can say something quite a bit stronger than that: practical ends have as their *raison d'être* the realization or implementation of basic ends; their "normativity" derives from their connection to basic ends.

But what does this involve? How do we get from basic ends to practical ends? Many pictures of the relationship are possible. One idea: p is a reason-giving end just in case it is a first-order end (i.e., p is a practical end iff p is a basic end). A better idea (which I'll defend in due course) is the Ends-Reasons Principle.

ENDS- REASONS PRINCIPLE (ERP)
A basic end is a practical end if and only if it is *actionable*.

Here is an informal definition of actionability (precisified in Sect. 6):

Definition 8 An end p is *actionable* with respect an information-state I iff for some available action α, I entails that α realizes p.

Ends aren't actionable per se, but *relative to an information-state*.[20] For example, while saving as many lives as possible (all ten) is plausibly a basic end, it's not a

---

[20] The careful reader may notice a formal resemblance between actionability and choosability (discussed in Footnote 18). To recap, K&M suggest that $\Box \phi$ is true only when ϕ is choosable, where ϕ is choosable roughly if ϕ is actionable in the sense of Definition 8. But as the discussion in Footnote 18 makes clear, K&M's choosability requirement does not have anything to do with the properties of the deontic selection function. Nor, then, does it have anything to do with our implicit project here—i.e., developing a normatively plausible, but formally precise, method for constructing a deontic ordering on possibilities from a

practical end, relative to the relevant information. By stipulation, the available actions are: block A, block B, and block neither. For each such action, the relevant information either rules it out that or leaves it open whether it realizes the end of saving everyone. So that end is not actionable. The claim that practical ends must be actionable amounts to the claim that "reasoning" (in perhaps a very loose sense of that word) about the implementation of basic ends properly makes use of all and only those ends *we know how to implement*, by lights of the relevant information.

Why endorse ERP? This is plausible on a picture of practical reasoning on which the considerations that can bear on reasoning about the realization and implementation of basic ends are properly restricted to those considerations that it is known, by lights of the relevant information, how to implement. In fact, something very much like this picture has been defended by in a recent paper by Hawthorne and Stanley (2008). Hawthorne and Stanley give a series of arguments in favor of the following epistemic constraint on treating a consideration as a reason.

REASON- KNOWLEDGE PRINCIPLE (RKP)
"Where one's choice is p-dependent,[21] it is appropriate to treat the proposition that p as a reason for acting iff you know that p." (p. 578)

Assuming RKP, it is not difficult to derive ERP (if we fudge a bit and identify knowledge and relevant information). ($\Rightarrow$) Suppose p is both a basic and practical end, and that some available action α realizes p. Then it is appropriate to treat the proposition that α realizes p as a reason for performing α. By RKP, it is known that performing α realizes p. So, by Definition 8, p is actionable. ($\Leftarrow$) In the other direction, suppose p is an actionable basic end. Then, for some available action α, our knowledge settles that doing α realizes p. So, by RKP, it is appropriate to treat the proposition that α realizes p as a reason for acting. Not just any action, of course: since p is a basic end, it is appropriate to treat the proposition that α realizes p as a reason *for performing α*.

So, practical ends—the sorts of things that it is appropriate to take as counting for or against an action in practical reasoning—vary with the relevant information.[22] Indeed, they vary in a systematic way, if we suppose, as seems plausible, that the set of basic ends is closed under logical weakening (so that if a basic end p entails q, q

---

Footnote 20 continued

body of goals/ends. To simplify the discussion of Footnote 18 somewhat, K&M's choosability partitions the space of *outcomes*, while our actionability partitions the space of *goals/ends* (i.e., the sorts of considerations that determine deontic orderings on possibilities, hence which *determine which outcomes are deontically preferred*). Requiring obligatory outcomes to be choosable tell us *nothing* about what properties a goal/end p must have to contribute to some p-possibility's relative deontic goodness. (If, however, the discussion here is on track, requiring that a goal/end p be actionable to contribute to some p-possibility's relative deontic goodness will tell us much about which outcomes are obligatory.)

To be clear, the account developed in this section is inspired by K&M's discussion in their Footnote 18. Theoretically, however, the formal similarity between actionability (of an end) and choosability (of an outcome) is incidental; actionability occupies a wholly different role in my account than choosability occupies in K&M's. I'm grateful to an anonymous reviewer for pressing me here.

[21] We will not worry about p-dependence, since all the considerations considered here are p-dependent. Nor will we worry about the sense of "appropriate" that RKP and the rest of our discussion invokes. My sketch requires only whatever detail is sufficient to guide the semantics in Sect. 6; we don't need to fuss about the precise understanding of concepts of propriety, knowledge, and so on.

[22] Basic ends may too (cf. our discussion of rule-case reasoning), just not in the way practical ends do.

is also a basic end). On that supposition, the *strongest actionable weakening* of any basic end will always count as a practical end.

Definition 9 If *p* is a basic end, *q* is a *maximally strong actionable weakening* of *p* iff (i) *q* is actionable, (ii) *p* entails *q*, (iii) ¬∃*r* : *p* entails *r*, *r* entails *q*, and *q* does not entail *r*.

How do these sorts of normative-cum-decision-theoretic considerations apply to the miner case? Their *semantic* import is not immediately obvious; we will develop it in the next section. But we can at least say the following. If it is a basic end that all ten miners are saved, it is a basic end that at least nine miners are saved (since the latter is strictly weaker than the former). The former end is not actionable relative to the relevant information. *But the latter is*. So the latter, but not the former, is a practical end. So, even if the miners are in A, it is not appropriate to treat the proposition that blocking A realizes the basic end of saving all ten as a reason to block A (since this is not a practical end). But it is appropriate to treat the proposition that blocking neither realizes the basic end of saving nine as a reason to block neither, since the relevant information entails that blocking neither realizes that end.

Notice also that *enriching the relevant information-state* tends to make it *easier for a basic end to be actionable*, hence, easier for it to be appropriate to treat the fact that an available action realizes that end as a reason for doing that action. Both of the above basic ends are actionable from the vantage of an information-state bearing the information that the miners are in shaft A. From the vantage of such an enriched information-state, then, it is appropriate to treat the proposition that blocking A will realize the basic end of saving all ten as a reason to block A.

This, basically, is the picture of practical reasoning we shall design our semantics for modals and indicatives around. It is plausible, but implementation in a semantics for modals and indicatives introduces technical wrinkles. We tackle these below.

6 Writing information-dependence into the semantics

The prior section furnished us with a distinction between basic and practical ends, as well as a fairly concrete account of the relationship between them. The remaining work of this paper is to translate these notions into a semantics for modals and indicatives that meets the three desiderata outlined in Sect. 4.

On the semantics developed here, the domain of quantification for weak deontic necessity modals is determined by an ordering on possibilities, constructed from a set of basic (first-order) ends, together with a set of secondary ends (that first-order ends be actionable with respect to relevant information). It is, in other words, determined relative to a set of practical ends. While the set of basic ends relevant for evaluating a weak deontic necessity modal is, more or less, *stable* modulo context-internal changes in the relevant information,[23] the set of practical ends relevant for evaluating such a

---

[23] In a small class of cases—*anankastic conditionals*, specifically—our semantics will predict (correctly; cf. von Fintel and Iatridou 2005) that entertaining an indicative's antecedent may shift the set of basic ends relevant to evaluating its consequent. I simply note this prediction here; I will not discuss it further.

modal is highly *unstable* modulo context-internal changes in the relevant information (e.g., suppositions, making certain possibilities salient for the sake of reasoning about their consequences, and so on).[24]

On our account, this sort of information-dependence—one which is *peculiar to weak deontic necessity modals*, rather than a reflex of a Stability violation for the deontic selection function—is what generates the requisite variation in the quantificational domains of weak deontic necessity modals. Let me emphasize that, while this is not the sort of account on offer from K&M, it should be interpreted as a friendly amendment to their account. Indeed, they might be inclined to embrace it (with the proviso that *their* deontic selection function be understood, not as selecting out best or ideal worlds *sans phrase*, but rather as selecting the out the worlds that are best relative to a set of practical ends, in roughly the way sketched in this section).

6.1 Ends and orderings

We will assume that, at a context c:

– The basic ends at c determine a unique deontic ordering on possibilities for c, hence also determine a unique *deontic selection function* for c.
– The practical ends at c determine a unique practical ordering on possibilities for c, hence also determine a unique *practical selection function* for c.

In this section, we will state a precise, if somewhat idealized, general recipe for getting from sets of ends to orderings and selection functions.

Let us first formalize the intended relationship between the set of basic ends at a context c and the set of practical ends at c.

Definition 10 At a context c and possibility i: if $\mathcal{B}_{c,i}$ designates the set of basic ends at $\langle c, i \rangle$, then the *set of practical ends at $\langle c, i \rangle$* (notation: $\mathcal{P}_{c,i}$) is the set of those $p \in \mathcal{B}_{c,i}$ such that p is actionable with respect to $F_c(i)$.

Letting $\mathcal{A}_{c,i}$ designate the actions available at $\langle c, i \rangle$, Definition 10 basically amounts to:

$$\mathcal{P}_{c,i} = \{\llbracket \phi \rrbracket^c \in \mathcal{B}_{c,i} : \exists \alpha \in \mathcal{A}_{c,i} : F_c(i) \subseteq \llbracket \alpha \text{ is done} \supset \phi \rrbracket^c\}$$

This gives us a recipe for getting a set of practical ends from a set of basic ends at a context. But how do we get from here to rankings on possibilities and selection functions? What we really need here is a notion of *comparative better-ness* of possibilities relative a set of propositions. A crude such notion—one that, nevertheless, has been influential in semantic treatments of natural language deontic modals—is due to Kratzer (1981). Relative to a set of propositions X (what Kratzer terms an "ordering source"), a possibility i is at least as good as another possibility j just in case, for any $p \in X$, if j satisfies p, then i does too.

---

[24] A relevant decision-theoretic precedent is the work on conditional utility of Weirich (1980).

Definition 11 If X is an ordering source and i and j are possibilities, i is *at least as X-good* as j (notation: $i \preceq_X j$) iff $\{p \in X \mid j \in p\} \subseteq \{p \in X \mid i \in p\}$.

This notion of comparative better-ness has its virtues; for instance, it allows the resulting ordering to be *partial*, rather than total (unlike a notion of comparative better-ness which ranked possibilities according to, e.g., the *quantity* of propositions in the ordering source they satisfied).

Proof Suppose that $X = \{p, q\}$; that $p \cap q = \emptyset$; and that $i \in p$ and $j \in q$. Then, since $p \cap q = \emptyset, i \notin q$ and $j \notin p$. So $\{p' \in X \mid j \in p'\} \not\subseteq \{p' \in X \mid i \in p'\}$ and $\{p' \in X \mid i \in p'\} \not\subseteq \{p' \in X \mid j \in p'\}$. So, by Definition 11, $i \not\preceq_X j$ and $j \not\preceq_X i$. □

This is nice if one thinks, as many value theorists do, that (i) incompatible ends (represented as incompatible propositions in an ordering source) are possible, (ii) incompatible ends tend to induce *incomparabilities* (i.e., gaps) in the comparative better-ness rankings they determine (cf. Swanson, forthcoming).

Kratzer's notion of comparative better-ness does have its difficulties. As with a counting notion of comparative better-ness, the Definition 11 notion makes something like an *Equal Footing* presupposition for ends: satisfying any end in X counts just as much in a possibility's favor (with respect to $\preceq_X$) as satisfying any other end in X. So, Equal Footing means idealizing away, e.g., any sort of hierarchy of importance among basic ends. For our purposes, the correctness of this idealization (and others) is immaterial: what matters is that we have a reasonably concrete, even if extremely idealized, picture of how a body of ends might determine an ordering on possibilities. As we will see, an idealized model might nevertheless still offer an illuminating way of implementing the picture of practical reasoning developed in the prior section.

What of selection functions? We will retain our earlier understanding of selection functions, updating it slightly to reflect the role of ordering sources in determining orderings on possibilities. Selection functions, as before, select the best worlds from a domain of possibilities, with respect to an ordering on possibilities. Since orderings are determined by ordering sources, selection functions must be relativized to ordering sources, in addition to the usual parameters.

Definition 12 The *selection function* relative to a context-world pair $\langle c, i \rangle$ and function g from context-world pairs to ordering sources (notation: $\sigma_{c,i,g}$) is a function from a domain p into the $\preceq_{g_{c,i}}$-minimal possibilities in p.

Note that, on this definition, selection functions meet a Realism constraint. Namely, for all c, i, g: $\sigma_{c,i,g}(F_c(i)) \subseteq F_c(i)$. We will also suppose they meet a Definedness constraint. Namely, for all c, i, g: $\sigma_{c,i,g}(F_c(i))$ is defined and non-empty. (For the original definitions of these constraints, see Sect. 2.)

Note that Definition 12 straightaway entails Stability *for an arbitrary selection function*.

Proof Let $\sigma_{c,i,g}$ be a selection function. Suppose $j \in \sigma_{c,i,g}(p)$. By Definition 12, j is $\preceq_{g_{c,i}}$-minimal in p. But then, by Definition 11, if $p' \subseteq p$ and $j \in p'$, j is $\preceq_{g_{c,i}}$-minimal in p'. So, by Definition 12, $j \in \sigma_{c,i,g}(p')$. □

Since we are assuming Stability, our solution to the miner puzzle *cannot* in any way rest on Stability violations in *any* selection function (deontic or otherwise). What is best, relative to a set of ends, is totally domain-independent. It is, instead, the *ends themselves*—practical ends, specifically—that are domain-dependent. We will make this a bit more precise, and finally cash it out in a semantics, in the next section.

6.2 Weak necessity

von Fintel and Iatridou (2008), drawing on an account of *ought* due originally to Sloman (1970), suggest that:

> [W]hen we use [(8)], what is conveyed is that there are several ways of going to Ashfield but that by some measure, Route 2 is the best:
> (8) To go to Ashfield, you ought to use Route 2.
> … What makes [weak necessity modals] weaker semantically is that… strong necessity modals say that the prejacent is true in all of the favored worlds, while weak necessity modals say that the prejacent is true in all of the very best (by some additional measure) among the favored worlds. In the terms of the Kratzerian framework, we suggested that weak necessity modals are in general sensitive to (at least) two ordering sources (pp. 118–119).

The hypothesis, in short (and informally): "saying that to go to Ashfield you ought to take Route 2, because it's the most scenic way, is the same as saying that to go to Ashfield in the most scenic way, you have to take Route 2" (2008, p. 137). Although they present some suggestive linguistic evidence for this hypothesis (cf. Footnote 19), discussing it is a bit beyond the scope of this paper; we will just need to rely on the hypothesis' substantive, normative plausibility (which is, nevertheless, considerable).[25] Here, we will just show how (i) making use of (a suitably modified version of) this hypothesis, while (ii) understanding the secondary ordering source in the miner context in a certain way yields a solution to the miner puzzle satisfying all three of our desiderata.

Let me first situate the miner case with respect to von Fintel and Iatridou's hypothesis about weak deontic necessity. I claim the miner context belongs to a certain class of contexts—call them *deliberative* (or, more dramatically, *risky*) contexts—in which the salient secondary considerations are higher-order: they bear on which first-order considerations are allowed to determine what we *should* do (i.e., allowed to determine the domain of weak deontic necessity modals). Here is a substantive, normative claim about the salient secondary considerations at the miner context: what we *should* do about the miners is a function of what considerations it is appropriate for us, given the

---

[25] Criticisms of von Fintel and Iatridou's proposal exist, but none, so far as I know, target the claim that weak necessity modals make use of primary and secondary ordering sources. Swanson, forthcoming, e.g., objects to the treatment of weak necessity modals as universal quantifiers, partly on the grounds that it validates agglomeration for weak necessity modals—i.e., predicts that *should* $\phi \wedge$ *should* ψ entails *should*$(\phi \wedge \psi)$. But this is completely independent of whether weak necessity modals make use of primary and secondary ordering sources. Whatever the correct semantics for weak necessity modals, so long as they make use of primary and secondary ordering sources, the proposal here will be adaptable to it.

relevant information, to treat as reasons for or against certain courses of action. At the miner context, then, secondary considerations function to determine what propositions it is appropriate to treat as reasons for or against certain courses of action. So, by the Ends-Reasons Principle (see Sect. 5.2), the salient secondary considerations demand actionability.

That seems plausible enough. But how to write it down in the semantics? On von Fintel and Iatridou's account, weak deontic necessity modals are sensitive to a set of *designated goals* (analogous to first-order ends) as well as a set of *secondary ends*. But sensitive how? One possibility is to have strong deontic necessity modals quantify over possibilities that are best with respect to first-order ends, while having weak deontic necessity modals quantify over possibilities that are best with respect to *both* first-order and secondary ends. Letting $\mathcal{B}$ and $\mathcal{S}$ designate, respectively, functions from context-world pairs into sets of first-order and secondary ends, this would mean endorsing the following pair of truth-conditions.

Definition 13 $[\![must\ \phi]\!]^{c,i} = 1$ iff $\sigma_{c,i,\mathcal{B}}(F_c(i)) \subseteq [\![\phi]\!]^c$

Definition 14 $[\![should\ \phi]\!]^{c,i} = 1$ iff $\sigma_{c,i,\mathcal{S}}(\sigma_{c,i,\mathcal{B}}(F_c(i))) \subseteq [\![\phi]\!]^c$

This, however, cannot be the way to formalize the sort of information-sensitivity I have claimed that weak deontic necessity modals exhibit in the miner case. For, notice that, if the body of first-order ends—$\mathcal{B}_{c,i}$—contains the proposition that we maximize the amount of life saved, then every possibility in $\sigma_{c,i,\mathcal{B}}(F_c(i))$ will tend to satisfy the proposition that we maximize the amount of life saved. Further contracting this domain, in the manner of Definition 14, cannot change this fact; we will thus tend to predict, *prima facie* incorrectly, that we should maximize the amount of life saved (hence, that we should block at least one shaft).

Information-sensitivity, then, is not here a matter of further narrowing a domain determined by a fixed body of first-order ends. Here, it is a matter of selecting out those first-order ends that are actionable (i.e., *selecting out the practical ends* from a set of first-order ends, in the manner of Definition 10), and determining a quantificational domain using the ordering on possibilities that the practical ends determine. Equivalently, in certain cases (e.g., the miner case), secondary ends can *coarsen the ordering* that determines the domain of weak deontic necessity modals, by removing non-actionable propositions from the ordering source.

Here, then, is a better idea for weak necessity modals:

Definition 15 $[\![should\ \phi]\!]^{c,i} = 1$ iff $\sigma_{c,i,\mathcal{B}\oplus\mathcal{S}}(F_c(i)) \subseteq [\![\phi]\!]^c$

The $\oplus$ operation on $\mathcal{B}$ and $\mathcal{S}$ is meant to express some sort of *merging* operation on these parameters. If c designates the miner context, what does merging $\mathcal{B}$ and $\mathcal{S}$ amount to at c? With our picture of information-sensitive practical reasoning (and our account of how this affects the ordering source utilized by weak necessity modals) in the background, this is actually a rather easy question to answer: for any possibility i, $(\mathcal{B}\oplus\mathcal{S})_{c,i}$ should be the set of propositions in $\mathcal{B}_{c,i}$ that *meet the constraints* specified

by $\mathcal{S}_{c,i}$ (in particular, being actionable at $\langle c, i \rangle$).[26] That's to say:

$$\forall i : (\mathcal{B} \oplus \mathcal{S})_{c,i} = \mathcal{P}_{c,i} \ \text{ (as defined in Definition 10)}$$

6.3 A solution (finally!)

This little claim, together with assumption that the correct semantics for indicatives is the Shifty Semantics (see Definition 6), represents our Official Solution to the miner puzzle. How does it do with respect to our desiderata?

Does it have a theory-external rationale? Yes. It grows out of three things: (i) an independently motivated account of how basic considerations of value and goodness properly bear on practical reasoning (codified in the Ends-Reasons Principle); (ii) an independently motivated semantics for weak necessity modals, which makes their quantificational domains sensitive to both primary and secondary ordering sources; (iii) a substantive claim that the miner context is a deliberative (or risky) context (in that the function of the relevant secondary ordering source is to select first-order ends that it is appropriate for us, given the relevant information, to treat as reasons for or against performing the available actions). All of (i)–(iii) are plausible, in contrast to K&M's dubiously motivated rejection of Stability.

Does it explain consistency by explaining truth? Yes. Let c be the miner context, i a possibility, and suppose that the body of basic or first-order ends at $\langle c, i \rangle$, $\mathcal{B}_{c,i}$, contains the proposition that we maximize the amount of life saved and the body of secondary ends at $\langle c, i \rangle$, $\mathcal{S}_{c,i}$, contains an actionability constraint on elements of $\mathcal{B}_{c,i}$. We will show that the sentences in (3)—hence, the sentences in (1)—are true at $\langle c, i \rangle$. (More formal versions of these explanations can be found in the footnotes.)

Why is $\neg should(block\_A \lor block\_B)$ true at $\langle c, i \rangle$? Informally, this is because the merged ordering source used to interpret this sentence does not contain the proposition that we maximize the amount of life saved; it is removed because it is not actionable with respect to the relevant information. So possibilities where we save nine (by doing nothing) are ranked just as highly, according to $(\mathcal{B} \oplus \mathcal{S})_c$, as those where we save ten (by blocking the shaft in which the miners are located).[27, 28]

---

[26] More difficult is to define a *general* merging operation holding between *arbitrary* sets of basic and secondary ends that yields the right results, not just for the miner case, but also for cases like (8). Now, compositionality does not, as such, require there be any such operation, so perhaps we could do without one. Although I don't think we ultimately have to, it's actually controversial what the merged ordering source is supposed to look like even in cases like (8) (see von Fintel and Iatridou 2008, p. 138). Giving a general definition for $\oplus$ would require settling this controversy. That is beyond the scope of this paper.

[27] More formally: by Definition 15, $[\![should(block\_A \lor block\_B)]\!]^{c,i} = 1$ iff $\sigma_{c,i,\mathcal{B}\oplus\mathcal{S}}(F_c(i)) \subseteq [\![block\_A \lor block\_B]\!]^c$. But $(\mathcal{B} \oplus \mathcal{S})_{c,i}$ does not contain the proposition that we maximize the amount of life saved; rather it contains the strongest actionable weakening of this proposition (namely, that we save at least nine). So, relative to $\preceq_{(\mathcal{B}\oplus\mathcal{S})_{c,i}}$, possibilities in $F_c(i)$ where we do nothing (and nine are saved) are ranked as highly as those where we block the correct shaft (and all ten are saved). So $\sigma_{c,i,\mathcal{B}\oplus\mathcal{S}}(F_c(i)) \not\subseteq [\![block\_A \lor block\_B]\!]^c$. Hence, $[\![should(block\_A \lor block\_B)]\!]^{c,i} = 0$ and $[\![\neg should(block\_A \lor block\_B)]\!]^{c,i} = 1$.

[28] Notice further that predicting the truth of $should\neg(block\_A \lor block\_B)$, if desired, can be secured by using a slightly modified version of the merged ordering source—one that treats non-actionable propositions

Why are the indicatives (*if in_A*)(*should block_A*) and (*if in_B*)(*should block_B*) true at $\langle c, i \rangle$? We'll consider the first; what goes for it goes for the other. Informally, according to the Shifty Semantics (Definition 6), (*if in_A*)(*should block_A*) is true iff (*should block_A*) is true throughout, and relative to, a domain bearing the information that the miners are in A; entertaining the antecedent shifts the information that is relevant for evaluating the consequent. But, relative to such a domain, the proposition that we maximize the amount of life saved is actionable. And so, relative to the ordering determined the ordering source for such a domain, possibilities where we do nothing are ranked less highly than those where we block shaft A (thereby saving everyone). So, indeed, (*should block_A*) is true throughout, and relative to, a domain bearing the information that the miners are in A. By parallel reasoning, (*should block_B*) is true throughout, and relative to, a domain bearing the information that the miners are in B. So both (*if in_A*)(*should block_A*) and (*if in_B*)(*should block_B*) are true at $\langle c, i \rangle$.[29]

Additionally, the semantics clearly offers the *right kind* of explanation of truth (as detailed in our discussion of Desideratum 2 (Sect. 4)). For it explains:

- How the computation of truth-values of the claims in (1) and (3) is sensitive to facts about preferred outcomes.
- How context-dependent limitations on information can affect the interpretation of descriptions of that agent's obligations.
- How conditionally introduced suppositions can exercise a peculiar effect on the interpretation of linguistic constructions expressing conditional obligation.

Does it explain the weak / strong asymmetry? Yes. We've made this point already, but it bears repeating: the information-sensitivity exploited to explain the truth of the sentences in (3) is *due entirely* to a feature that is peculiar to weak necessity modals, namely, their sensitivity to secondary concerns (in the miner case, concerns of actionability). This is a substantial empirical improvement over a solution like K&M's, which trades on rejecting Stability for *the* deontic selection function. A K&M-style solution has trouble explaining the felt inconsistency of the sentences in (5). This is because it introduces information-sensitivity, in the form of Stability violations, into

---

Footnote 28 continued

as strictly undesirable, rather than regarding them neutrally. We can define this modified ordering source $(\mathcal{B} \oplus \mathcal{S})^*_{c,i}$ to be the result of intersecting every element of $(\mathcal{B} \oplus \mathcal{S})_{c,i}$ with the *negation* of every *non-actionable* element of $\mathcal{B}_{c,i}$:

$$(\mathcal{B} \oplus \mathcal{S})^*_{c,i} = \{p : \exists q \in (\mathcal{B} \oplus \mathcal{S})_{c,i} : \exists r \subset q \in \mathcal{B}_{c,i} \notin (\mathcal{B} \oplus \mathcal{S})_{c,i} : p = q \cap \bar{r}\}$$

$(\mathcal{B} \oplus \mathcal{S})^*_{c,i}$ will, in other words, contain the proposition that we save at least nine, but do not save all ten, miners. Actions that save exactly nine miners will be strictly preferred, according to this ordering source, to all other actions. And so it will follow that *should*¬(*block_A* ∨ *block_B*). (I'm grateful to Kai von Fintel for pressing me on this.)

[29] More formally: by Definition 6, $[\![(if\ in\_A)(should\ block\_A)]\!]^{c,i} = 1$ iff $F_{c+in\_A}(i) \models_{c+in\_A} should\ block\_A$, iff, by Definition 4, $F_{c+in\_A}(i) \subseteq [\![should\ block\_A]\!]^{c+in\_A}$, iff, by Definition 15, for any $j \in F_{c+in\_A}(i) : \sigma_{c+in\_A, j, \mathcal{B}\oplus\mathcal{S}}(F_{c+in\_A}(j)) \subseteq [\![block\_A]\!]^{c+in\_A}$. But notice that $(\mathcal{B}\oplus\mathcal{S})_{c+in\_A, j}$ contains the proposition that we maximize the amount of life saved. That is because, relative to $F_{c+in\_A}(j)$, this proposition is actionable. So, relative to $\preceq_{(\mathcal{B}\oplus\mathcal{S})_{c+in\_A, j}}$, possibilities in $F_{c+in\_A}(j)$ where we do nothing (and nine are saved) are ranked *lower* than those where we block A (since, for any possibility k in $F_{c+in\_A}(j)$, blocking A maximizes the amount of life saved at k). So $\sigma_{c+in\_A, j, \mathcal{B}\oplus\mathcal{S}}(F_{c+in\_A}(j)) \subseteq [\![block\_A]\!]^{c+in\_A}$.

the deontic selection function—the same selection function utilized by *all* deontic necessity modals, weak and strong alike.

6.4 Performativity and strong modality

There is a residual puzzle about strong deontic necessity modals, which I shall address briefly here. If strong deontic necessity modals are, as we've suggested, sensitive only to the primary ordering source, and the primary ordering source contains the proposition that we maximize the amount of life saved, then we would expect (9) to be true at the miner context.

(9) $must(block\_A \lor block\_B)$

But, of course, that is problematic. For one, (9) certainly seems to be false. Moreover, if $\neg should(block\_A \lor block\_B)$ is true at the miner context, then, since *must* entails *should*, (9) *must be false* there.

My reply, in two parts: the relevant reading of (9)—paraphrasing, roughly, as *to save everyone, we must either block A or B*—is literally true at the miner context. However, even on the indicated reading:

1. (9) is not utterable at the miner context (which also explains our unwillingness to assent to it).
2. (9) does *not* entail (the salient reading of) $should(block\_A \lor block\_B)$.

Unacceptability as infelicity. Regarding (i), it is usually recognized that *strong* necessity modals carry some sort of *performative* force (see esp. Ninan 2005; Portner 2007). What this means, very roughly, is that (a) utterances of strong necessity modals conventionally express proposals to get the addressee to intend to make their prejacents true, (b) accepting a strong necessity modal (including one addressed to oneself) means accepting its prejacent as something to be intended by the addressee. Such facts are, for independent reasons, thought to explain the impossibility of expressing strong deontic necessities while thinking the addressee will fail to satisfy them (since one cannot, rationally, attempt to create an intention while simultaneously expecting that that intention will go unfulfilled or fail to be formed at all, as in examples 10 and 11), as well as the possibility of expressing weak deontic necessities while still thinking that the addressee will fail to satisfy them, as in examples (12) and (13) (see esp. Ninan (2005)).

(10) #You must go to church, although you won't.

(11) #I must go to church, although I won't.

(12) You should go to church, although you won't.

(13) I should go to church, although I won't.

Analogously, one cannot express or accept a strong deontic necessity while rejecting or disavowing its conventional performative upshot. Applying this to (9), the fact that, to save everyone, we must either block A or B does not mean that (9) can be felicitously uttered at the envisioned context. In light of the performativity of (9), such an

utterance is strange, on the supposition that one does not wish to enjoin the realization of (*block_A* ∨ *block_B*). Similarly, given the performativity of *must*, we cannot happily accept (9), if we do not accept its prejacent as something to be intended. And, of course, we do not. Such facts do not vitiate our semantic proposal; the performativity of strong modals is a *non-semantic* fact about their conventional performative force.

Here is a somewhat more fleshed out version of this explanation. The suggestion, roughly, is that while weak necessity modals "directly" exploit secondary considerations at the level of semantic interpretation (via use of secondary ordering sources), strong necessity modals *"indirectly" exploit* secondary considerations at the level of presupposition (or, more neutrally, on the level of conditions on felicitous/appropriate use/utterance by a rational agent). Strong necessity modals conventionally express a special kind of performative force—they conventionally express proposals that their addressees come to intend their prejacents (a fact of which competent speakers are aware). Often an agent cannot *rationally* attempt to generate an intention in her addressee that ϕ, even when ϕ is required for the realization of something she regards as desirable. In such cases uttering *must* ϕ will be generally infelicitous. In some such cases, this is because she believes that her efforts will be futile, as in (10) and (11). In others, as I'm suggesting is the case with (9), this is because she thinks ϕ *inadvisable*. In the case of (9), the prejacent (*block_A* ∨ *block_B*) is inadvisable because it is not knowable that realizing (*block_A* ∨ *block_B*) will lead to the realization of things she regards as desirable (and may well, given what she knows, lead to the realization of things she regards as undesirable); these things are non-actionable with respect to her limited information. Since she regards (*block_A* ∨ *block_B*) as inadvisable (which results from her bringing secondary considerations to bear on her practical reasoning), she cannot rationally attempt to generate an intention in her addressee with content (*block_A* ∨ *block_B*). So it is not felicitous for her to utter (9) at the envisioned context.

Note that while secondary considerations are certainly exploited in this sort of reasoning, they are not exploited *within the semantics* for strong necessity modals. Rather, they figure in an account of the "mechanics" of practical reasoning by a rational agent under conditions of subjective uncertainty. In short, (9) is neither rationally utterable nor acceptable in a context where its prejacent is regarded as inadvisable. And the prejacent is regarded as inadvisable because the salient primary consideration whose realization requires the prejacent to be true (namely, saving everyone) is non-actionable with respect to the agent's limited information.

When must implies should. Regarding (ii), the salient reading of *should*(*block_A* ∨ *block_B*) is, on our proposal, paraphrasable as *to do well with respect to our actionable primary goals, we must block A or B*. But the fact that, to do well with respect to our primary goals (sans phrase), we have to block A or B, does not, for reasons we have already seen, imply that to do well with respect to our *actionable* primary goals, we have to block A or B. So (9) fails to entail (the salient reading of) *should*(*block_A* ∨ *block_B*).

This may seem like an unacceptably high cost. But the cost, I think, is not nearly so high as it seems. When interpreted with respect to a *fixed* deontic ordering on possibilities, *must* does still entail *should*. So, for instance, in any case where $\mathcal{B}_{c,i} = (\mathcal{B} \oplus \mathcal{S})_{c,i}$, it will turn out that $[\![must\ \phi]\!]^{c,i} = 1$ implies $[\![should\ \phi]\!]^{c,i} = 1$, since the deontic

ordering is constant: $\preceq_{\mathcal{B}_{c,i}} = \preceq_{(\mathcal{B} \oplus \mathcal{S})_{c,i}}$.[30] The only entailment we are actually relinquishing is the entailment from (14) to (15):

(14) To do well with respect to our primary goals, we must block *A* or *B*

(15) To do well with respect to our actionable primary goals, we must block *A* or *B*

But, of course, no one should accept *this* entailment, given the suggested definition of actionability.

To put the point a bit differently, it would not be generally regarded as an objection to a semantics for *must* that *must* ϕ might fail to entail *must* ϕ when the relevant deontic ordering is allowed to vary between episodes of evaluation. Of course it can be the case that all the best worlds with respect to a body of goals G are ϕ-possibilities without it being the case that all the best worlds with respect to a distinct body of goals G' are ϕ-possibilities. For instance, the only way to get to Harlem via public transportation might be via the "A" Train, even when the best way to get to Harlem, regardless of mode of transportation, is by taxi. But this is just what my semantics for weak and strong deontic necessity modals claims is happening in the case of (9). The content of (9) is paraphrased as (14). The content of $should(block\_A \vee block\_B)$ is paraphrased as (15). (14) uncontroversially fails to entail (15).

The critic of this paper's proposed analysis therefore owes us an account of why this failure of entailment should worry us more than the failure of *must* ϕ to entail *must* ϕ when the relevant deontic ordering is allowed to vary between episodes of evaluation. One idea might be that an agent who accepts *must* ϕ is always licensed to infer *should* ϕ. But the acceptability of such an inference can be explained without relinquishing our account. By the pragmatic account sketched above, a rational agent who accepts *must* ϕ necessarily regards ϕ as advisable (where ϕ is advisable roughly just if ϕ stands in the same relation to the agent's information and her primary and secondary concerns as it stands when *should* ϕ is the case). But then a rational agent will not regard ϕ as advisable unless she is disposed to accept *should* ϕ. Since such an agent is rational, an agent who accepts *must* ϕ is always licensed to infer *should* ϕ.

---

[30] Since, on the suggested semantics, *must* and *should* both express $\forall$, it will turn out that *should* entails *must* as well when the deontic ordering is held fixed (since both will quantify universally over the *same domain*). Now, that seems implausible: strong and weak necessities are intuitively non-equivalent, *even with respect to a fixed deontic ordering*. But it is well known that treating deontic necessity as universal quantification is a useful theoretical idealization (since it is approximately correct for most cases). To gesture at one alternative, Eric Swanson's "ordering supervaluationism" (Swanson, forthcoming), on which evaluating *must* involves supervaluating over sets of deontically comparable possibilities, while evaluating *should* does not—is especially well-suited to accounting for the asymmetry of this entailment, even with respect to a fixed deontic ordering on possibilities. It does not require giving up the primary/secondary ordering source machinery that is central to our account.

References

Burnyeat, M. (1980). Aristotle on learning to be good. In A. O. Rorty (Ed.), *Essays on Aristotle's ethics*, Chap. 5. University of California Press.

Cariani, F., Kaufmann, S., & Schwager, M. (ms). Deliberative modality under epistemic uncertainty. Unpublished.

Dreier, J. (2009). Relativism (and expressivism) and the problem of disagreement. *Philosophical Perspectives, 23*, 79–110.

Gillies, A. S. (2010). Iffiness. *Semantics and Pragmatics, 3*, 1–42.

Hansson, S. O. (2005). Preference logic. In D. M. Gabbay & F. Guenthner (Eds.), *Handbook of philosophical logic* (2nd ed., pp. 319–393). Berlin: Springer.

Hawthorne, J., & Stanley, J. (2008). Knowledge and action. *The Journal of Philosophy, 105*, 571–590.

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

Kratzer, A. (1981). The notional category of modality. In H. Eikmeyer & H. Rieser (Eds.), *Words, worlds, and contexts* (pp. 38–74). Berlin: De Gruyter.

Kratzer, A. (1991). Conditionals. In A. von Stechow & D. Wunderlich (Eds.), *Semantics: An international handbook of contemporary research* (pp. 651–656). Berlin: De Gruyter.

Lewis, D. (1973). *Counterfactuals*. Malden: Basil Blackwell.

Lewis, D. (1975). Adverbs of quantification. In Keenan E. (Ed.), *Formal semantics of natural language* (pp. 3–15). Cambridge: Cambridge University Press.

MacFarlane, J. (2011). Epistemic modals are assessment-sensitive. In B. Weatherson & A. Egan (Eds.), *Epistemic modality*. Oxford: Oxford University Press.

Ninan, D. (2005). Two puzzles about deontic necessity. In J. Gajewski, V. Hacquard, B. Nickel, & S. Yalcin (Eds.), New work on modality, Vol. 51 of MIT working papers in linguistics. Cambridge: MITWPL.

Parfit, D. (ms). What we together do.

Portner, P. (2007). Imperatives and modals. *Natural Language Semantics, 15*, 351–383.

Sen, A. K. (1969). Quasi-transitivity, rational choice and collective decisions. *The Review of Economic Studies, 36*, 381–393.

Silk, A. (ms). Evidence-sensitivity in deontic modals. Unpublished.

Sloman, A. (1970). 'Ought' and 'better'. Mind, *LXXIX*, 385–394.

Stalnaker, R. (1968). A theory of conditionals. In N. Rescher (Ed.), *Studies in logical theory (American Philosophical Quarterly Monograph Series 2).* (pp. 98–112). Oxford: Blackwell.

Stephenson, T. (2008). Judge dependence, epistemic modals, and predicates of personal taste. *Linguistics and Philosophy, 30*, 487–525.

Swanson, E. (2008). Modality in language. *Philosophy Compass, 3*, 1193–1207.

Swanson, E. (forthcoming). On the treatment of incomparability in ordering semantics and premise semantics. *Journal of Philosophical Logic*. Available at http://www.springerlink.com/content/p84253x306652274/.

von Fintel, K., & Gillies, A. S. (2008). CIA leaks. *The Philosophical Review, 117*, 77–98.

von Fintel, K., & Iatridou, S. (2005). What to do if you want to go to Harlem: Anankastic conditionals and related matters. Ms., MIT.

von Fintel, K., & Iatridou, S. (2008). How to say *ought* in foreign: The composition of weak necessity modals. In J. Guèron & J. Lecarme (Eds.), *Time and modality*. Springer Netherlands.

Weirich, P. (1980). Conditional utility and its place in decision theory. *The Journal of Philosophy, 77*, 702–715.

Wiggins, D. (1975). Deliberation and practical reason. *Proceedings of the Aristotelian Society, 76*, 29–51.

Yalcin, S. (2007). Epistemic modals. *Mind, 116*, 983–1026.