

Deontic Conditionals: Weak and Strong^{*}

Alex Silk
University of Birmingham
a.silk@bham.ac.uk

Draft of July 2015

Abstract

Discussions of the Miners Puzzle have prompted renewed interest in how deontic modals can be sensitive to relevant information. Linguistic phenomena involving information-sensitivity have been used to motivate general claims about the interpretation of modals and conditionals. These generalizations have been premature. Theorists have focused on a limited range of cases. Nearly all examples use weak necessity modals ('ought', 'should'), and little if any attention is paid to how context affects speakers' judgments. This leaves existing theories ill-equipped to capture further data (presented in this paper) involving possibility modals and strong necessity modals in a range of discourse contexts. Indeed I argue that puzzles concerning information-sensitivity arise because of features unique to weak necessity modals. We can explain the data concerning information-sensitivity without revising our understanding of modals and conditionals in general. All that is needed is an independently motivated semantics for weak necessity modals. Considering the broader spectrum of examples elucidates the variety of expressive resources at our disposal for coordinating our actions and expectations in conversation, deliberation, and planning.

^{*}Initial draft first posted Winter 2012. A previous version of this paper was published in *Proceedings of the 19th International Congress of Linguists*, Département de Linguistique de l'Université de Genève, 2013. Thanks to Fabrizio Cariani, Jan Dowell, Josh Knobe, Bob Stalnaker, Eric Swanson, Steve Yablo, Malte Willer, and audiences at the University of Michigan, MIT, the 19th International Congress of Linguists, and the 2013 Eastern APA for discussion and judgments.

Discussions of the Miners Puzzle have prompted renewed interest in how deontic modal claims can be sensitive to relevant information. The following is a description of the Miners Case and a standard account of the data that it provides.

Ten miners are trapped either in shaft *A* or in shaft *B*, but we do not know which. Flood waters threaten to flood the shafts. We have enough sandbags to block one shaft, but not both. If we block one shaft, all the water will go into the other shaft, killing any miners inside it. If we block neither shaft, both shafts will fill halfway with water, and just one miner, the lowest in the shaft, will be killed. (KOLODNY & MACFARLANE 2010: 115–116)¹

Since we don't know which shaft the miners are in, and the consequences will be disastrous if we choose the wrong shaft, (1) seems true.

(1) We should block neither shaft.

The conditionals in (2)–(3) are also natural to accept. After all, given that the miners are in Shaft *A* (/Shaft *B*), blocking Shaft *A* (/Shaft *B*) will save all the miners.

(2) If the miners are in Shaft *A*, we should block Shaft *A*.

(3) If the miners are in Shaft *B*, we should block Shaft *B*.

(2)–(3), along with (4), which we know, seem to entail (5).

(4) Either the miners are in Shaft *A* or the miners are in Shaft *B*.

(5) Either we should block Shaft *A* or we should block Shaft *B*.

But (5) contradicts (1). Hence the miners' puzzle.

Many recent authors have argued that cases like the Miners Case pose a serious challenge for the classic Kratzerian semantics for modals and conditionals.² Even those who have eschewed revising the classic semantics still use these cases to motivate claims about the interpretation of modals and conditionals in general.³ These reactions have been premature. Discussions of information-sensitivity in deontic

¹The case was originally discussed in PARFIT 1988, credited to Donald Regan. See, e.g., JACKSON 1991 for similar cases.

²See, e.g., KOLODNY & MACFARLANE 2010, LASSITER 2011, CARIANI ET AL. 2013, CHARLOW 2013, MACFARLANE 2014, SILK 2014b.

³See DOWELL 2012, VON FINTEL 2012.

modals and conditionals have focused on a narrow range of cases. Nearly all examples are with weak necessity modals, like ‘ought’ and ‘should’, and little if any attention is paid to how context affects speakers’ judgments. Proceeding in this way runs the risk of drawing hasty generalizations and producing a theory that is ill-equipped to capture a broader spectrum of examples.

Indeed, I will argue that what makes the Miners Puzzle a puzzle has nothing to do with deontic modals and conditionals *as such*. Puzzles concerning information-sensitivity arise because of features unique to weak necessity modals like ‘should’. We can explain the data concerning information-sensitivity without revising our understanding of modals and conditionals in general. All that is needed is an independently motivated semantics for weak necessity modals. Considering a broader spectrum of examples — e.g., involving possibility modals and strong necessity modals, in various discourse contexts — elucidates the variety of expressive resources at our disposal for coordinating our actions and expectations in conversation, deliberation, and planning.

The primary aims of this paper are, first, to introduce a broader range of data concerning information-sensitivity in deontic modals and conditionals; second, to highlight previously underappreciated contrasts and sensitivities to context in deontic ‘should’ and ‘must’ claims, both unembedded and in conditionals; and, third, to develop a more nuanced account of deontic modals that is sensitive to these phenomena. Though some of the broader data I will provide have been previously noted in the literature, there has been little investigation of how to theorize about them. The project isn’t to argue that no other theory can get the data right. It is to motivate one way of capturing intuitions driving previous responses to the Miners Puzzle that improves in empirical coverage and independent motivation. I welcome the development of alternative ways of systematizing the data with which the account offered here may be compared.

Roadmap: §1 presents new data concerning the Miners Case. §2 argues that these data raise challenges for existing accounts of information-sensitivity in deontic modals and conditionals. §3 offers an independently motivated semantics for weak necessity modals like ‘should’. §4 shows how this semantics can help make the correct predictions about the broader range of data concerning the Miners Case.

1 Information-sensitivity with ‘must’ and ‘may’

We need more data on information-sensitivity in deontic modals and conditionals. This section provides a range of broader examples. These examples reveal contrasts

in information-sensitivity with ‘should’, on the one hand, and ‘must’ and ‘may’, on the other.

I want to flag that reactions to some of the examples may be vague, for some speakers, and can vary given subtle changes in context. One might worry about building a theory on the basis of such data. If our intuitions about the original data with ‘should’ are sufficiently clear, why not take them as our starting point in theorizing? However, there is reason to think that much of the original data, and what to make of it, is less clear than initially seemed. We need more detailed specifications of context and greater attention to how subtle variations in context affect judgments. This holds not just for the additional examples provided throughout this paper, but for the original examples with ‘should’ as well. The proper theoretical import of the variety of examples, and of how context affects judgments, are important parts of what needs to be explained. The positive account in §§3–4 will, I hope, crystalize the informal reactions described in this section.⁴

(I use ‘should’ in examples, as opposed to ‘ought’, simply because I find it more natural. I focus on the strong necessity modal ‘must’ and possibility modal ‘may’ because, like ‘ought’ and ‘should’, they are typically used “subjectively” in the sense of LYONS 1977 — they typically present the speaker as endorsing the considerations that would verify the modal claim. Non-speaker-endorsing uses (more common with, e.g., ‘have to’, ‘supposed to’, ‘can’) introduce complications that would be distracting to our discussion here. For this reason it is important to consider the examples with the specific modals being used. That said, for speakers who find deontic ‘must’ generally awkward, ‘gotta’ or ‘have got to’ are arguably most similar to ‘must’ in the relevant respects.⁵)

Intuitively, the ‘should’ conditionals in (2)–(3) say what is best on a condition: given that the miners are in Shaft A (/Shaft B), our blocking Shaft A (/Shaft B) is the expectably best action. (2)–(3) don’t impose obligations on us conditional on how the world happens to be, unbeknownst to us. By contrast, (6)–(7) do seem to impose such an obligation.⁶

(6) ?If the miners are in Shaft A, we must block Shaft A.

(7) ?If the miners are in Shaft B, we must block Shaft B.

⁴Thanks to an anonymous referee for pressing me to clarify these methodological issues.

⁵See SWANSON 2012, SILK 2012b, 2014a, 2015 for recent discussion.

⁶I use ‘?’ to indicate that using the marked item is dispreferred. Thus ‘?’ marks a weaker infelicity than ‘#’. CHARLOW (2013: 2305–2306) reports similar judgments about the ‘must’ conditionals, both from himself and from informants. I discuss Charlow’s account in §2.

This is likely part of why many speakers find using ‘must’ in the conditionals to be dispreferred to using ‘should’. Informally, using ‘must’ in the conditionals seems to be more emphatic than is warranted. However, the relative felicity of ‘should’ and ‘must’ can vary depending on the context. Consider a context in which it is particularly urgent that no miners be lost. Suppose we are members of a rescue team called upon to save the miners. We have a longstanding record of never losing anyone on the job, and we intend to keep that record intact. After all, people’s lives are at stake! Someone might die! Getting into this frame of mind can improve the felicity of the ‘must’ conditionals in (6)–(7).

It is true that some speakers still have difficulty not hearing the ‘must’ conditionals as anomalous. What is important is that *if* one can at least imagine entering into a sort of urgent context in which the ‘must’ conditionals are acceptable, (1), and even (8), no longer seems true. Alice’s response in (9) is marked.

- (8) We may block neither shaft.
- (9) *Alice:* People’s lives are at stake! If they’re in A, we must block A. And if they’re in B, we must block B.
Bert: So what should we do?
Alice: #We should (/may) block neither shaft.

The ‘must’ conditionals in (6)–(7), along with (4), appear to be jointly inconsistent with (1), and even with (8).

Surprisingly, similar points appear to hold for ‘may’. Consider the following variant of the Miners Case — call it ‘Miners*’. As before, ten miners are trapped either in Shaft A or Shaft B, though we don’t know which, and we only have enough sandbags to block one shaft. But now suppose that Shaft A is connected to another shaft, Shaft C, in such a way that blocking Shaft A or Shaft C will save all the miners if they are in Shaft A but will kill all the miners if they are in Shaft B. Shaft B is likewise connected to another shaft, Shaft D, in such a way that blocking Shaft B or Shaft D will save all the miners if they are in Shaft B but will kill all the miners if they are in Shaft A. Many find using ‘may’ in (10)–(13) to be anomalous, where this anomalousness seems similar in character to that from using ‘must’ in (6)–(7).

- (10) ?If the miners are in Shaft A, we may block Shaft A.
- (11) ?If the miners are in Shaft B, we may block Shaft B.
- (12) ?If the miners are in Shaft A, we may block Shaft C.
- (13) ?If the miners are in Shaft B, we may block Shaft D.

It isn't counterintuitive that part of what sounds odd about using 'may' in (10)–(13) is that it grants us a permission to block such-and-such shaft conditional on how the world happens to be, independent of our knowledge, similarly to how using 'must' in (6)–(7) appears to impose such an obligation. The 'should' conditionals in (2)–(3) are acceptable in a way that the 'may' conditionals in (10)–(11) are not. Deontic 'should' does not imply deontic 'may'. (Hereafter I will focus specifically on the 'may' conditionals in (10)–(11) to emphasize the comparisons with the 'should' and 'must' conditionals in (2)–(3) and (6)–(7), respectively.)

What is most important for present purposes, however, is that *if* the 'may' conditionals in (10)–(11) are accepted, (1) no longer seems true. It is hard to hear the 'may' conditionals as acceptable except in a context in which (14) is also accepted, as brought out in (15)–(16).

- (14) We must block one of the shafts.
- (15) *Alice*: If they're in A, we may block A. And if they're in B, we may block B.
Bert: So what should I tell the team? Should I flip a coin to decide which shaft to block?
Alice: #What are you talking about? We can't block either shaft!
- (16) *Alice*: If they're in A, we may block A. And if they're in B, we may block B.
Bert: But we don't know which shaft they're in. Why not hedge our bets and guarantee saving nine of the miners?
Alice: Because people's lives are at stake! We must block one of the shafts!

The Miners Puzzle is no puzzle at all when expressed with 'must' or 'may'.

To recap, our first piece of new data concerns intuitions of relative felicity: using 'must' in the deontic conditionals is dispreferred to using 'should' in the original Miners Case, and using 'may' in the deontic conditionals is more dispreferred in the variant Miners* than using 'should' in the deontic conditionals is in the original Miners Case. Importantly, however, these relative felicities are not absolute, but can vary given certain shifts in context (more on which below). Our discussion highlights the methodological importance of considering a variety of contextualized cases when soliciting speaker judgments. The second piece of data concerns intuitions about natural language entailments: the 'should' conditionals in (2)–(3), unlike the 'must' conditionals in (6)–(7) or the 'may' conditionals in (10)–(11) (in Miners*), are jointly consistent with (1) and (4).

Preliminary polling suggests that the intuitions recorded in this section are sufficiently robust to warrant beginning theorizing about them; they aren't idiosyncratic (see also n. 6). But it is worth noting that there is also a sense in which they are ex-

pected in light of independent, established data concerning deontic modals in root clauses. There is a robust body of evidence, much of which from descriptive linguistics, that deontic ‘must’ and ‘may’ are typically used performatively in root clauses: utterances of deontic ‘Must ϕ ’ typically issue an imperative, and utterances of deontic ‘May ϕ ’ typically grant a permission. By contrast, even if deontic ‘Should ϕ ’ can be used to perform an imperatival speech act in certain contexts, its directive force is typically much weaker.⁷ Seen against the backdrop of these broader data, the data in this section are perhaps less surprising. They may be understood as an extension of the data concerning unembedded uses to the case of another type of matrix environment or main clause, namely, the consequent of a conditional.⁸ This connection with performativity, of course, doesn’t constitute an explanation of the present data. Indeed it locates an important aspect of what needs to be explained — namely, why information in the antecedent of a ‘should’ conditional, unlike a ‘must’ or ‘may’ conditional, is hypothetically taken on board in interpreting and assessing the force of the modalized consequent clause (more on this in §2). Also, it isn’t implausible that we should seek to derive the contrasts in performativity and directive force from more basic features of the conventional meanings of the modals and general pragmatic principles. My point in introducing the data about performativity is simply to highlight how the data introduced in this section coheres with independent data on deontic modals. Observing this connection may be helpful for readers who find themselves questioning some of the judgments I have recorded. In any case, I invite readers with alternative judgments to join me in investigating how we might theorize about the dialect of English characterized here.

2 Previous treatments

The examples in §1 suggest that deontic ‘must’ and ‘may’ aren’t subject to the same sorts of apparent modus ponens violations as deontic ‘should’, and that contextual

⁷Paul McNamara (1990: 156) characterizes the phenomena well: “To say that one ought to take a certain option is merely to provide a nudge in that direction. Its typical uses are to offer guidance, a word to the wise (“counsel of wisdom”), to recommend, advise or prescribe a course of action... In contrast, to say that one must take a certain option is to be quite forceful. Its typical uses are to command, decree, enact, exhort, entreat, require, regulate, legislate, delegate, or warn. Its directive force is quite strong.” See also, a.o., WERTHEIMER 1972, WILLIAMS 1981, COATES 1983, PALMER 1990, MYHILL 1996, HUDDLESTON & PULLUM 2002, WERNER 2003, NINAN 2005, SWANSON 2008, PORTNER 2009, SILK 2012b.

⁸For related discussion in the case of conditional imperatives see STARR 2010, KAUFMANN & SCHWAGER 2011; cf. the treatment of conditional questions in ISAACS & RAWLINS 2008.

factors can affect the relative felicity of deontic ‘should’, ‘must’, and ‘may’ conditionals. Previous treatments of information-sensitivity in the literature leave these points unexplained. For space purposes I will focus on the accounts in KOLODNY & MACFARLANE 2010 (as a representative of a revisionary, information-sensitive semantics), VON FINTEL 2012 (as a representative of the classic semantics), and CHARLOW 2013 (as the only response in the literature that addresses the relevance of the contrast between weak and strong necessity modals in the Miners Case). (Additional accounts are briefly considered in note 12.)

Start with the semantics in KOLODNY & MACFARLANE 2010. On this semantics deontic modals are interpreted with respect to a body of information i , represented by a set of possible worlds, and a deontic selection function d that selects a set of best worlds from this domain: ‘Should ϕ ’ is true iff ‘ ϕ ’ is true at all worlds $w' \in d(i)$.⁹ What makes the semantics “information sensitive” in an interesting, revisionary sense is that the selection function is sensitive to which worlds are eligible. Even if a world is selected as “best” relative to some domain, it might not be selected relative to a subset of that set of which it is a member.¹⁰ This property of the selection function helps capture how a deontic modal in the consequent of a conditional can be interpreted with respect to the information carried by the antecedent: roughly, ‘If ψ , should ϕ ’ is true iff ‘ ϕ ’ is true at all worlds $w' \in d(i \cap \psi)$.¹¹ This semantics predicts that (1)–(4) are jointly consistent. Though the best worlds relative to our information state, which is unsettled as to the location of the miners, are worlds in which we block neither shaft, the best worlds relative to our information state strengthened with the information that the miners are in, say, Shaft A are worlds in which we block Shaft A.

Kolodny and MacFarlane’s account is, at best, incomplete: it is silent on the contrasts between weak and strong necessity modals and their relevance for the Miners Puzzle, and it makes no provisions for how shifts in context affect the relative felicities of different deontic modals. (There is, after all, only so much one can do in

⁹Kolodny & MacFarlane treat d as supplied by the context of utterance, though they treat the informational parameter i as an added parameter of the index (circumstance of evaluation). This feature of their view concerning the internal structure of points of evaluation won’t be relevant in what follows.

¹⁰See CHARLOW 2013 and SILK 2014b for further discussion.

¹¹“Roughly” because Kolodny & MacFarlane ultimately revise this semantics, and treat the ‘should’ conditional as true iff ‘ ψ ’ is true at all worlds $w' \in d(i' \cap \psi)$, for every maximal subset i' of i throughout which ‘ ϕ ’ is true. This complication won’t be relevant in what follows.

I treat ‘ ϕ ’, ‘ ψ ’, etc. as schematic letters to be replaced with declarative sentences. For convenience I sometimes refer to the proposition expressed by ‘ ϕ ’ by dropping the single quotes — e.g., using ‘ $i \cap \phi$ ’ as short for ‘ $i \cap \llbracket \phi \rrbracket^c$ ’, where $\llbracket \phi \rrbracket^c$ is the set of worlds in which ‘ ϕ ’ is true.

a single paper.) But given certain natural assumptions, the mechanisms by which they respond to the Miners Puzzle overgenerate and predict the felicity of the ‘may’ and ‘must’ conditionals and their joint consistency with (1). Let ‘ blX ’ denote the proposition that we block Shaft X, and let ‘ inX ’ denote the proposition that the miners are in Shaft X. First, if we treat ‘may’ as expressing existential quantification over the set of deontically best worlds, as would be expected, then the ‘may’ conditionals in (10)–(11) are predicted to be unambiguously true and consistent with (1). If all worlds in $d(i \cap inA)$ are blA -worlds, then some worlds in $d(i \cap inA)$ are blA -worlds. Second, insofar as ‘should’ is already given a semantics of necessity, expressing universal quantification over the set of best worlds, the semantics, at least as it stands, has no mechanism for predicting the contrasting discourse properties of ‘should’ and ‘must’. But however the semantics might be refined to make ‘must’ logically stronger than ‘should’, it is unclear why the ‘must’ conditionals in (6)–(7) would not also be acceptable or consistent with (1) insofar as the deontic selection function for a deontic modal embedded in the consequent of a conditional is given as argument the strengthened information state — i.e., the initial information state updated with the information carried by the antecedent.

Kolodny and MacFarlane might respond by interpreting ‘must’ and ‘may’ with respect to a different deontic selection function, one that isn’t information-sensitive. But it is false that (e.g.) ‘must’ claims never depend on relevant information or what we know. (Analogous points hold for ‘may’.) Using ‘must’ becomes significantly more natural if we settle definitively that we won’t learn which shaft the miners are in, as in (17), or if we do learn which shaft the miners are in, as in (18) (more on this point below).

- (17) *Alice*: Well, it looks like we’re never going to be able to find out which shaft the miners are in.
Bert: So what should I tell the team?
Alice: That we must block neither shaft.
- (18) *Alice*: I figured it out; the miners are in Shaft A. Tell the team: we must block Shaft A!

One might weaken the proposed revision and allow that (e.g.) ‘must’ can sometimes be interpreted with respect to an information-sensitive deontic selection function. But this still fails to provide any detail on how the different selection functions are determined and what unifies them. In our examples, the values and norms that bear on what we ought to do are, intuitively, the same as those that bear on what we must and may do. Other things equal, I think we should seek an account that

captures this point.

To be clear, I am not claiming that Kolodny and MacFarlane's basic account cannot be revised or supplemented to capture the additional data from §1. My point is simply that their basic account leaves much to be explained. Independent resources, in addition to the revisionary features of their semantics, will be required to capture the relevant context-sensitivities and differences among deontic conditionals. This makes pressing whether introducing these revisionary features provides the most helpful starting point for theorizing about our broader data.¹²

VON FINTEL 2012 proposes no revisions to the standard semantics in responding to the Miners Puzzle. In his attempt to save the standard semantics, von Fintel denies that the 'should' conditionals in (2)–(3) are literally true. He treats them, instead, as elliptical for conditionals like (19)–(20).

(19) If we learn that the miners are in Shaft A, we should block Shaft A.

(20) If we learn that the miners are in Shaft B, we should block Shaft B.

There are a number of concerns with this sort of response (see CARR 2012 for further discussion). What I want to emphasize here is that it overgenerates. There is a clear contrast in acceptability between 'must' and 'may' conditionals with 'if ψ ' and 'if we learn that ψ ' as their antecedents. Even in a context in which (6)–(7) and (10)–(11) are anomalous, the following conditionals are perfectly felicitous.

(21) If we learn that the miners are in Shaft A, we must block Shaft A.

(22) If we learn that the miners are in Shaft B, we must block Shaft B.

(23) If we learn that the miners are in Shaft A, we may block Shaft A.

¹²Analogous challenges hold for the information-sensitive semantics in LASSITER 2011, CARIANI ET AL. 2013, and SILK 2014b. The semantics in CARIANI ET AL. 2013 and SILK 2014b have the advantage of integrating Kolodny and MacFarlane's discussions into a standard ordering semantics framework. The extra structure in their accounts may provide greater resources for capturing some of the contrasts discussed in §1. But, again, more work will need to be done. Most pressingly, the posited mechanisms for interpreting modals in conditional consequents with respect to the information carried by their antecedents will need to be reined in for the cases of 'may' and 'must', and done so in such a way that leaves room for the relevant sensitivities to context. LASSITER 2011 treats deontic modals as making claims about expected values. Roughly, 'Must (/May) ϕ ' is true iff the expected value of ϕ exceeds some high (/low) threshold value; and 'Should ϕ ' is true iff the expected value of ϕ is greater than the expected value of any contextually relevant alternative ψ . Deontic conditionals are treated as claims about conditional expected values. This semantics positively predicts that the 'must' and 'may' conditionals in (6)–(7) and (10)–(11) are true and consistent with (1). It also makes no provisions for how shifts in context affect the relative felicity of 'should' and 'must'.

(24) If we learn that the miners are in Shaft B, we may block Shaft B.

It would be odd, to put it mildly, if the antecedents of deontic ‘should’ conditionals were reinterpreted in the proposed way but the antecedents of deontic ‘must’ and ‘may’ conditionals were not.

CHARLOW 2013 is alone in the literature on information-sensitivity in recognizing that ‘should’ and ‘must’ differ in their sensitivities to relevant information. On Charlow’s semantics, ‘must’ quantifies over the best worlds in an ordering that ranks worlds in terms of their realization of certain basic values — e.g., the number of lives saved. This basic ordering ranks worlds where we block the correct shaft as best, worlds where we do nothing as next best, and worlds where we block the incorrect shaft as worst. Information-sensitivity with ‘should’ is captured by coarsening this ordering. In the coarsened ordering that figures in the interpretation of ‘should’, worlds cannot be ranked differently unless they differ with respect to the realization of our actionable values, i.e. those we know how to realize. Since we don’t know which shaft the miners are in, the value of saving all ten isn’t actionable in this sense. Hence worlds where we block neither shaft (and save nine miners) are ranked along with worlds where we block the correct shaft (and save all ten) as best. This predicts that ‘We should block either shaft’ is false.

Charlow’s sensitivity to differences among necessity modals is commendable. But I have two concerns about the specifics of his proposal.¹³ First, Charlow’s account incorrectly predicts that the truth of a ‘must’ claim never depends on relevant information or what we know. As we saw in (17), if we settle that we will not learn the miners’ location, using ‘must’ becomes acceptable, even preferred. But Charlow’s semantics predicts that (25)

(25) We must block neither shaft.

is false since it treats ‘must’ as only being interpreted with respect to the relevant basic values — here, the number of lives saved.

Second, Charlow’s semantics predicts that the ‘must’ and ‘may’ conditionals in (6)–(7) and (10)–(11) are all true. (Charlow doesn’t consider variants with ‘may’ conditionals, but I assume he would treat ‘may’ as the dual of ‘must’.) But even

¹³In addition, though Charlow’s account predicts that ‘We should block either shaft’ is false, it fails to predict that (1) is true. This is because the set of most highly ranked worlds in the coarsened ordering still includes worlds where we block the correct shaft. (The ordering relevant for ‘should’ merely coarsens the ordering relevant for ‘must’; if a world is best relative to the refined ordering, it can’t fail to be best relative to the coarsened ordering.) For discussion see CHARLOW 2013: n.26, LASSITER 2011: 136n.6. Given my purposes I won’t press this worry here.

if there are contexts in which these conditionals are true, they are not true in all contexts (§1). Drawing on his remarks about a similar issue (2013: 30–32), Charlow might reply that the ‘must’ and ‘may’ conditionals are unassertable even if true. He might say that in light of the performative properties of ‘must’ and ‘may’, utterances of these conditionals impose an obligation or grant a permission, respectively, conditional on how the world happens to be; but since blocking one of the shafts is inadvisable, given our limited information, performance of such speech acts would be generally infelicitous. But, first, one might wonder why similar considerations wouldn’t predict the ‘should’ conditionals to be similarly infelicitous. Insofar as ‘should’ claims still have directive force, albeit weaker than that of ‘must’, why shouldn’t utterances of the ‘should’ conditionals (say) recommend, advise, or suggest a course of action conditional on how the world happens to be (*cf.* n. 7)? Second, we need an account of precisely what features of context affect the relative felicity of ‘should’, ‘must’, and ‘may’. Again, more work needs to be done.

The data from §1 raise a challenge for existing accounts of information-sensitivity. The challenge is to revise or supplement their account of how the ‘should’ conditionals are true and jointly consistent with (1) in such a way that (*a*) doesn’t similarly predict the truth of the ‘must’ or ‘may’ conditionals and their consistency with (1), and (*b*) captures the ways in which our evaluations of various types of deontic modals and conditionals are sensitive to context. The worries raised in this section can be seen as introducing constraints on a response to this challenge. I haven’t argued that existing accounts of information-sensitivity cannot be made to work. My main aim is simply to motivate a positive alternative strategy for capturing phenomena concerning information-sensitivity. It is to this alternative that I now turn.

3 Semantics for ‘should’, ‘must’, and ‘may’

The considerations in §2 make appealing the following hypothesis: An account of the Miners Puzzle won’t lie simply in a denial of modus ponens or in general considerations about the semantics of deontic modals and conditionals. Weak necessity modals hold the key. In the remainder of the paper I will argue that an account of information-sensitivity and the Miners Puzzle falls out of a standard general semantic framework for modals and conditionals, and an independently motivated semantics for weak necessity modals. In this section I will briefly motivate and describe this account of the contrast between ‘should’, on the one hand, and ‘must’ and ‘may’, on the other, an account which I have developed in greater detail elsewhere (SILK 2012a,b). In the next section I will show how this account makes the correct predic-

tions about our data. The treatment of the Miners Puzzle developed here provides an attractive alternative to the sorts of accounts considered in §2.

There is a robust body of evidence supporting a distinction in strength among necessity modals.¹⁴ For instance, even holding the readings of the modals fixed, ‘Should ϕ ’ can be followed by ‘Must ϕ ’, but not vice versa, as reflected in (26).

- (26) a. I should help the poor. In fact, I must.
b. I must help the poor. #In fact, I should.

Similarly, (27a) is consistent in a way that (27b) is not.

- (27) a. I should help the poor, but I don’t have to.
b. #I must help the poor, but it’s not as if I should.

There are also important conversational differences between ‘should’ and ‘must’. For instance, as we saw in §§1–2, the relative felicity of ‘should’ and ‘must’ depends on standing assumptions in the context. It is this feature of weak and strong necessity modals that I would like to focus on here (see RUBINSTEIN 2012, SILK 2012a,b for extensive recent discussion).¹⁵

Suppose I am considering whether to fight in the Resistance or take care of my ailing mother. I mention that the value of family, which supports my helping my mother over my fighting, is important, and you agree. But the issue is acknowledged to be complex, and it isn’t settled in the conversation whether there might be more important competing values. Sensitive to this, you may find it more appropriate to express your advice that I help my mother by using ‘should’ than by using ‘must’, as in (28).

- (28) *Me*: Family is very important.
You: I agree. You should (/?must) take care of your mother.

But if we settle that family is of primary importance, as in (29), it can become more natural to use ‘must’ and for us to accept that I have to help my mother.

- (29) *Me*: Family is most important — more important than country.

¹⁴See, e.g., SLOMAN 1970; HORN 1972, 1989; WERTHEIMER 1972; HARMAN 1975, 1993; LYONS 1977; WILLIAMS 1981; COATES 1983; MCNAMARA 1996; PALMER 1990, 2001; HUDDLESTON & PULUM 2002; VON FINTEL & IATRIDOU 2008; SILK 2012b. I will continue to use ‘should’ as my representative weak necessity modal, and ‘must’ as my representative strong necessity modal (§1).

¹⁵See WOISETSCHLAEGER 1977: ch. 5 and MCNAMARA 1990: ch. 3 for prescient discussion and related cases. See also VON STECHOW ET AL. 2006: 14–15, VON FINTEL & IATRIDOU 2008: 139–140.

You: I agree. You must take care of your mother.

This kind of case highlights a crucial difference between weak necessity modals like ‘should’ and strong necessity modals like ‘must’. In both (28) and (29) the value of family is accepted in the conversation. Where (28) and (29) differ — and what affects how you express your advice — is our assumptions about the status of this value vis-à-vis other potentially competing values. My having an actual obligation to help my mother depends on the value of family being more important (or at least not less important¹⁶) in my situation than any competing value. Using ‘must’ is preferred if it is settled in the conversation that this condition obtains. What is illuminating is that you can felicitously express your advice that I help my mother using ‘should’, advice which I may accept, even if it isn’t common ground that this precondition for my having a genuine obligation to help my mother is satisfied. Accepting your ‘should’ claim needn’t require us to presuppose that the value of family is more important than other potentially competing values.

This conversational contrast between ‘should’ and ‘must’ is perhaps even clearer in the case of epistemic readings of the modals. Suppose we are working on an art project, and I ask you where the colored pencils are. Normally you put them in the drawer with the crayons, but sometimes you accidentally put them on the shelf. In this scenario it is more appropriate for you to use ‘should’ in responding to my question, as in (30).

(30) *Me:* Do you know where the colored pencils are?

You: They should (/?must) be in the drawer with the crayons.

Suppose, alternatively, that we are looking for the colored pencils together, and you indicate that you have just seen something that leads you to conclude that they are in the drawer. Perhaps you noticed that they weren’t on the shelf, and this is the only other place you think they could be. In this scenario it is more natural for you to use ‘must’, as in (31).

(31) *Me:* Do you know where the colored pencils are?

You: They must (/?should) be in the drawer with the crayons.

How you express your attitude toward the proposition that the colored pencils are

¹⁶For present purposes I will bracket potential complications concerning incomparabilities and genuine dilemmas (irresolvable conflicts of obligations). Note that the semantics offered in the following sections allows dilemmas to be consistently expressed with ‘should’ but not with ‘must’; see SWANSON 2011, SILK 2013b: ch. 3, 2015, and references therein, for relevant discussion.

in the drawer depends on the (in)defeasibility of the relevant evidence. Its following from our knowledge (evidence, information) that the colored pencils are in the drawer depends on today's not being one of the atypical days when you accidentally put the colored pencils on the shelf. Using the strong necessity modal 'must' is preferred if, and only if, you know that conditions are indeed normal in this way. Parallel to the deontic case, what is illuminating is that you can use 'should' even if you aren't in a position to judge that they are. Accepting your 'should' claim doesn't require us to presuppose that your evidence is indefeasible.

The lessons I wish to draw from these examples are twofold. First, 'Should ϕ ' doesn't conventionally communicate that ϕ is in fact necessary (deontically, epistemically, etc.). We can accept your deontic 'should' claim in (28) without settling that family is the most important relevant value, and thus without accepting that I have an actual obligation to help my mother. And we can accept your epistemic 'should' claim in (30) without settling that conditions are normal in the relevant respects, and thus without accepting that our evidence (knowledge, information) actually entails that the colored pencils are in the drawer. It is typical — to a first approximation — to gloss deontic notions of necessity as concerning what is obligatory, and epistemic notions of necessity as concerning what follows from one's evidence (knowledge, information).¹⁷ In this sense accepting deontic 'Ought ϕ ' doesn't commit one to accepting that ϕ is deontically necessary, and accepting epistemic 'Ought ϕ ' doesn't commit one to accepting that ϕ is epistemically necessary. Second, whether 'should' or 'must' is preferred depends on context, in the sense of depending on whether certain preconditions for the truth of the necessity claim (in the above sense) are accepted. If they are accepted, 'must' is preferred. But even if they aren't, we can still use 'should'.

The remainder of this section describes one way of capturing these points.¹⁸ The core of the proposal is this: There is nothing special about the semantics of 'must' or 'may'. Strong necessity modals are to be given their usual semantics of necessity. 'Must/May ϕ ' says that ϕ is necessary/possible, and predicates the necessity/possibility of ϕ of the actual world. The apparent weakness of weak necessity modals derives from their bracketing the assumption that the necessity of ϕ holds in the actual world.

Let's start by examining the sorts of considerations that figure in the interpretation of modals. It is a commonplace that values, norms, preferences, expectations,

¹⁷See, e.g., LYONS 1977, KRATZER 1978, COATES 1983, PALMER 1990, 2001, SWEETSER 1990, VAN DER AUWERA & PLUNGIAN 1998, HUDDLESTON & PULLUM 2002, a.m.o.

¹⁸See RUBINSTEIN 2012 for an alternative take on the data; see SILK 2012a,b for critical discussion.

etc. often come with conditions under which they apply. If I want to go for a run, my desire needn't be that I go for a run, come what may. More plausibly it is that I go for a run given that it's sunny, that I didn't just eat a burrito, and so on. Our preferences are often conditional, preferences *for* certain circumstances. Similarly with moral norms. Suppose you promised Alice that you would help her move. A norm against breaking your promise might be something to the effect that you help Alice unless you made a conflicting promise to Bert, or keeping your promise would lead to some serious harm, and so on. Norms can thus be understood on the model of *conditional* imperatives, imperatives that enjoin an action or state of affairs given that certain circumstances obtain. This captures the intuitive idea that depending on the circumstances, only certain norms, etc. apply, or are "in force." Fixing terminology, I will call a conditional norm, etc. a *consideration*. For a consideration ϕ if C , C is the consideration's *condition*, and ϕ is the consideration's *premise*, or what the consideration *enjoins* given C . A consideration *applies* at a world w if its condition is satisfied at w . (Categorical considerations can be treated as conditional on the tautology, and hence apply at any world.)

These points about the considerations with respect to which modals are interpreted can be integrated into a standard premise semantic framework for modals. Following Angelika Kratzer (1977, 1981, 1991), I treat modals as receiving their intended reading (deontic, epistemic, etc.) from a contextually supplied set of premises.¹⁹ Since modals can occur in intensional contexts, it is standard to index premise sets to a world of evaluation. What Little Timmy's parents command might change from one world to the next. This motivates treating the meaning of 'what Timmy's parents prescribe' in (32) as a function that assigns to every possible world the set of propositions describing the house rules in that world.

(32) In light of what Little Timmy's parents prescribe, he must be in bed by eight.

Such functions from worlds to premise sets are what context supplies for the interpretation of modals. Following Kratzer, call these functions *conversational backgrounds* (written ' P '). Call the value of a conversational background at a world of evaluation a *premise set* (written ' P_w ').

Conversational backgrounds afford a natural way of representing the contents of bodies of considerations. Suppose we have a consideration which enjoins ϕ given that conditions C obtain. We can represent the content of this consideration with a conversational background P that assigns to every relevant C -world a premise

¹⁹Equivalently, a preorder (LEWIS 1981). See also VAN FRAASSEN 1973, LEWIS 1973, VELTMAN 1976.

set that includes ϕ . For example, we can represent the normative import of your promise to help Alice move (see above) with a deontic conversational background that assigns a premise set that includes the proposition that you help Alice to worlds in which you made this promise, you didn't make a certain promise to Bert, and so on. The premises in a premise set thus reflect what follows from a body of considerations — what is enjoined by a body of conditional norms, what is preferred in light of a body of conditional preferences, etc. — given the relevant circumstances in the evaluation world. An indexed premise set P_w represents the conclusions of the relevant considerations, given the facts in w .²⁰

With this background at hand let's turn to the semantics. I give 'must' and 'may' their usual semantics of necessity and possibility, respectively. Roughly put, 'Must ϕ ' says that the prejacent proposition ϕ follows from P_w — where P_w is the premise set that is the value of the given conversational background P at the evaluation world w (formal details shortly). Treating 'may' as the dual of 'must', 'May ϕ ' says that ϕ is compatible with P_w .²¹ What is important to note is that the truth of 'Must ϕ ' depends on the value of P at the world of evaluation. In order for 'Must ϕ ' to be true at w , given a conversational background P , the circumstances in w must be such that the premise set P_w entails ϕ . Assuming the evaluation world is the actual world $w_{@}$ — or at least a world which, for all that is presupposed in the context, might be the actual world — this predicts that accepting 'Must ϕ ' commits one to accepting that ϕ is a necessity (in the relevant sense) at $w_{@}$. (*Mutatis mutandis* for 'May ϕ ')

The same doesn't hold for 'should'. As we saw above, accepting 'Should ϕ ' doesn't require accepting that ϕ is in fact necessary. What makes weak necessity modals "weak," on my view, is that they bracket whether the necessity of the prejacent holds in the actual world (SILK 2012b). Unlike with 'must', in uttering 'Should ϕ ' one needn't assume that the actual circumstances are such that what the relevant considerations (norms) enjoin given those circumstances entails ϕ ; one needn't commit

²⁰On this formalization, information about the conditions of particular considerations is encoded in, but not itself recoverable from, a conversational background and its values across worlds. See my 2012a: 45–47, 2013a, 2014c for further developments and motivations for this way of understanding the standard premise semantic framework.

²¹For simplicity I assume that our premise sets are consistent, and I don't distinguish between Kratzer's modal bases and ordering sources. Dropping these assumptions, the simplified talk about what follows from P_w can be understood as short for talk about what follows from all maximally consistent subsets of $F_w \cup G_w$ that include F_w , where F_w is a modal base that describes some set of relevant background facts and G_w is an ordering source that represents the content of a relevant ideal at the evaluation world. This still makes the limit assumption (LEWIS 1973: 19–20) that ordering consistent subsets of $F_w \cup G_w$ that include F_w by set inclusion \subseteq yields a set of subsets that are \subseteq -maximal. For semantics without the limit assumption, see LEWIS 1973, KRATZER 1981, 1991, SWANSON 2011.

to being in a world w in which P_w entails ϕ . ‘Should’ expresses necessity only *on the supposition* that certain circumstances relevant to the necessity claim obtain. It needn’t be presupposed that these circumstances actually do obtain. The worlds in which the necessity of the prejacent is said to hold — the set of worlds w such that $\cap P_w \subseteq \llbracket \phi \rrbracket^c$ — needn’t be identified with the set of live possibilities in the conversation (the “context set (STALNAKER 1978),” the set of worlds compatible with what is mutually presupposed for the purposes of the conversation). (We will return to the question of what accepting ‘Should ϕ ’ *does* commit one to shortly.)²²

This feature of ‘should’ certainly isn’t the only distinguishing feature of weak necessity modals, nor does it mark the only sense of “strength” along which modals can differ. But, I will argue, it does carve out a crucial role for expressions of weak necessity in conversation, deliberation, and planning. This will help explain the conversational differences between ‘should’ and ‘must’ discussed above.

Weak necessity modals afford a means of coordinating on the implications of our values, norms, etc. without having to decisively settle how they apply and weigh against one another in particular circumstances. Return to (28)–(29). The relevant deontic conversational background P represents the content of the relevant norms and values (normative considerations), and encodes their relative priorities and under what conditions they apply. The indexed deontic premise set P_w represents what is ultimately enjoined by the normative considerations which apply in w given the circumstances in w . The deontic necessity of my helping my mother — my having an obligation to help her — depends on the value of family being more important in my situation than other potentially competing values (n. 16). Hence in order for my helping my mother to follow from P_w , it must be the case that the value of family takes priority in my situation in w . Accepting the ‘must’ claim ‘You must take care of your mother’ thus requires committing for the future course of the conversation to being in a world w in which this condition is satisfied and the applicable normative considerations P_w entail that I help my mother. But in (28), unlike in (29), after my assertion is accepted it still isn’t settled whether the value of family *does* take priority in my situation. So, were you to use ‘must’ you would imply that you are

²²As noted below, there are various ways of implementing these points in the formal semantics and pragmatics. But even at the present level of abstraction, this approach differs crucially from the other main approaches to weak necessity — e.g., comparative possibility/probability approaches and domain restriction approaches. For instance, the latter analyses maintain that accepting ‘Should ϕ ’ requires accepting that ϕ is a necessity (in some sense) at the actual world; it is just that the domain of ‘should’ is treated as a subset of the domain of ‘must’ (cf. VON FINTEL & IATRIDOU 2008, RUBINSTEIN 2012). The present approach rejects both of these claims (see SILK 2012a,b for detailed discussion). Thanks to an anonymous referee for encouraging me to clarify this.

foreclosing certain possibilities that I have left open. Unless you are in a position to do so, your using ‘must’ is dispreferred. By using ‘should’ you can leave open the possibility that the value of family might ultimately be outweighed or defeated. If I accept your ‘should’ claim, we can provisionally proceed as if my helping my mother is required without needing to settle that the value of family is more important than other competing values we accept or may come to accept.

There are various ways of implementing the above features of ‘should’, ‘must’, and ‘may’ in the formal semantics and pragmatics. However, it will be helpful for working through the examples from §1 to have specific analyses at hand. To fix ideas I adopt the following technical implementation (see SILK 2012b for developments and discussion of several static and dynamic alternatives).

There is nothing specially “strong” about strong necessity modals. They are just necessity modals with the usual general indicative presupposition that the worlds being talked about are included in the context set. One way of implementing this general indicative presupposition is to treat it as restricting the domain of the interpretation function to proper points of evaluation, i.e. to contexts c and worlds w such that $w \in c$. Applying this to a necessity modal and possibility modal yields the following semantics for ‘must’ and ‘may’, respectively:

Definition 1. $\llbracket \text{Must } \phi \rrbracket^c = \lambda w: w \in c. \cap P_w \subseteq \llbracket \phi \rrbracket^c$

Definition 2. $\llbracket \text{May } \phi \rrbracket^c = \lambda w: w \in c. \cap (P_w \cup \{\llbracket \phi \rrbracket^c\}) \neq \emptyset$

Uttering ‘Must/May ϕ ’ predicates the necessity/possibility of ϕ throughout (i.e., at each world in) the context set — no different from how uttering an atomic sentence ‘ ϕ ’ predicates ϕ throughout the context set. There is nothing special about the semantics of ‘must’ or ‘may’.

What distinguishes ‘should’ is that one can accept ‘Should ϕ ’ without having to presuppose that all the preconditions for the necessity of ϕ are satisfied, and without having to accept that ϕ is in fact necessary. One way of capturing this is to treat weak and strong necessity modals equivalently at the level of truth conditions, and to distinguish them in their presuppositions: weak necessity modals cancel the usual general indicative presupposition that the worlds being talked about are in the context set. This is reflected in Definition 3.

Definition 3. $\llbracket \text{Should } \phi \rrbracket^c = \lambda w. \cap P_w \subseteq \llbracket \phi \rrbracket^c$

What distinguishes ‘should’ from ‘must’, on this line, is that ‘should’ allows “improper” points of evaluation; ‘should’ lacks the presupposition that (all) the relevant worlds at which the prejacent is necessary are live possibilities. In uttering ‘Should ϕ ’

one makes a claim about the necessity of ϕ but fails to mark one's utterance as being about worlds that are candidates for actuality. Formalizing the analysis in this way raises the question of which worlds *are* being talked about in uses of 'should'. We will return to this in §4.2.

Conditionals are hard. There is considerable debate about how to capture the “epistemic flavor” and presuppositions of indicative conditionals generally. Adding issues concerning information-sensitivity in deontic modals, and a presuppositional treatment of the weak/strong necessity distinction, only makes matters worse. I will assume that something like the following intuitive idea is correct: just as ‘Must/May ϕ ’ is accepted iff ϕ is a necessity/possibility at each world in the context set, likewise ‘If ψ , must/may ϕ ’ is accepted iff ϕ is a necessity/possibility at each ψ -world in the context set. There are various ways of securing this, or something close to it, in the semantics and pragmatics — e.g., depending on one's broader views on restrictor vs. operator approaches to conditionals, the presence of covert modals in (some) explicitly modalized conditionals, the relation between truth conditions and acceptance conditions for conditionals, and the role of mood in (non-)indicative conditionals. For concreteness I will treat deontic conditionals as checking whether the modalized consequent clause is verified at the relevant worlds in which the antecedent is true, as reflected in Definitions 4–6, where R_w is the set of relevant worlds from w .²³

Definition 4. $\llbracket \text{If } \psi, \text{ must } \phi \rrbracket^c = \lambda w: w \in c. \{w': w' \in R_w \wedge w' \in \llbracket \psi \rrbracket^c\} \subseteq \{w': \bigcap P_{w'} \subseteq \llbracket \phi \rrbracket^c\}$

Definition 5. $\llbracket \text{If } \psi, \text{ may } \phi \rrbracket^c = \lambda w: w \in c. \{w': w' \in R_w \wedge w' \in \llbracket \psi \rrbracket^c\} \subseteq \{w': \bigcap (P_{w'} \cup \{\llbracket \phi \rrbracket^c\}) \neq \emptyset\}$

Definition 6. $\llbracket \text{If } \psi, \text{ should } \phi \rrbracket^c = \lambda w. \{w': w' \in R_w \wedge w' \in \llbracket \psi \rrbracket^c\} \subseteq \{w': \bigcap P_{w'} \subseteq \llbracket \phi \rrbracket^c\}$

For simplicity I will assume that $R_w = \{w\}$, i.e. that it is the evaluation world which is potentially relevant for assessing the deontic status of the modal's prejacent. Given our analyses for ‘must’, ‘may’, and ‘should’, this has the result that accepting deontic ‘If ψ , must/may ϕ ’ requires accepting that ϕ is deontically necessary/possible at every live ψ -world; ‘should’ conditionals, in contrast, place no such constraint. How to derive this compositionally will depend on difficult general questions concerning the syntax/semantics/pragmatics of (non-)indicative conditionals, weak necessity, mood, and interactions among them. Fortunately we can bracket many of these

²³This implementation coheres well, for instance, with double modal analyses of (some) overtly modalized conditionals. See, e.g., FRANK 1996, adapting Kratzer's (1981, 1991) classic restrictor analysis (*cf.* LEWIS 1973, 1975).

issues here. What is important for our purposes is just that the use of ‘should’, unlike ‘must’ or ‘may’, cancels the general indicative presupposition. In accepting a ‘should’ conditional, the set of relevant antecedent-worlds needn’t be identified with the set of antecedent-worlds in the context set.

4 Information-sensitivity in context

4.1 ‘Must’ and ‘may’

Let’s apply these analyses to the Miners Case and our data from §1. Suppose we are in a context — like our original context in the Miners Case — in which we take our obligations to depend on our epistemic state (knowledge, evidence, beliefs, available information, etc.). This dependence can be reflected in a body of conditional norms, norms which enjoin certain actions given what our epistemic state is like at the time. We can encode the contents of such norms in a simplified conversational background like P in (33), which assigns premise sets that include the proposition that we block neither shaft (blN) to worlds in which we don’t know which shaft the miners are in, and premise sets that include the proposition that we block Shaft X (blX) to worlds where we know the miners are in Shaft X (n. 20).²⁴ (AK , $A\bar{K}$, BK , $B\bar{K}$ can be treated as representatives of suitable equivalence classes of worlds characterized with respect to which shaft the miners are in and whether we know it.)

(33) *Evidence-dependent conversational background P:*

$$\begin{aligned} P_{AK} &= \{blA\} \\ P_{BK} &= \{blB\} \\ P_{A\bar{K}} &= P_{B\bar{K}} = \{blN\} \end{aligned}$$

²⁴Given our purposes I leave open whether the conversational backgrounds utilized in the main text are directly supplied by context, or are generated from more basic elements represented in the semantics. The fact that (e.g.) we have an obligation to block neither shaft in circumstances where we don’t know which shaft the miners are in can be derived from facts about the basic values at stake (e.g., saving miners’ lives), the available acts, and our epistemic state, as reflected in a theory of rational choice. It is an open question, however, whether these latter sorts of facts must be explicitly represented in the *semantics*, or conventional contents, of deontic modal sentences. (For responses to the Miners Puzzle that explicitly invoke this extra structure, see KATZ ET AL. 2012, CARIANI ET AL. 2013, CHARLOW 2013.) Given our aim of delivering the intuitively correct truth conditions in various contexts, it is sufficient to use conversational backgrounds that encode the eventual contents of the relevant (conditional) norms, and how our obligations do/don’t depend on the evidence. Doing so doesn’t trivialize our task of capturing the variety of data from §1. Thanks to an anonymous referee for pressing me on this issue. See SILK 2013a for general discussion.

(I am abstracting away from several complications concerning time. It wouldn't be misleading to treat the indices on P — and on the other conversational backgrounds to follow — as being restricted to time intervals $[t, t']$ beginning with the present point in our deliberation t , and culminating with the end of our deliberation t' , the time at which we must act or do nothing.)

With a context in which such a conversational background is called for in mind, reconsider the 'must' conditionals in (6)–(7).

- (6) If the miners are in Shaft A, we must block Shaft A.
- (7) If the miners are in Shaft B, we must block Shaft B.

In order for (6) to be accepted in the context, it must be the case that all worlds in the context set where the miners are in Shaft A verify the deontic necessity of our blocking Shaft A. So, accepting (6) would require removing world $A\bar{K}$ from the context set, since blA doesn't follow from $P_{A\bar{K}}$. But removing this world would be inapt. Insofar as we don't know which shaft the miners are in and don't want to presuppose that we will learn which shaft they are in, $A\bar{K}$ should remain a live possibility. This correctly predicts that (6) is rejected in our original context. (*Mutatis mutandis* for (7).)

However, suppose now that we are in a context like the one in (9), in which we take ourselves to have an obligation to block the shaft the miners are actually in. Such a context might call for a conversational background like P' in (34), which encodes the contents of norms that enjoin actions irrespective of our epistemic state.

- (34) *Evidence-independent conversational background P' :*

$$\begin{aligned} P'_{AK} &= P'_{A\bar{K}} = \{blA\} \\ P'_{BK} &= P'_{B\bar{K}} = \{blB\} \end{aligned}$$

The 'must' conditionals (6)–(7) are accepted in this sort of context: blA follows from P'_{AK} and $P'_{A\bar{K}}$, the values of P' at all live worlds where the miners are in Shaft A; and blB follows from P'_{BK} and $P'_{B\bar{K}}$, the values of P' at all live worlds where the miners are in Shaft B. But if (4) is also accepted, then (8) is false at all worlds in the context set and is thus rejected.

- (4) Either the miners are in Shaft A or the miners are in Shaft B.
- (8) We may block neither shaft.

For no world in the context set is the value of P' at that world compatible with the proposition that we block neither shaft.

Analogous points hold concerning the ‘may’ conditionals in (10)–(11).

- (10) If the miners are in Shaft A, we may block Shaft A.
 (11) If the miners are in Shaft B, we may block Shaft B.

Reconsider the variant of the Miners Case, Miners*, described in §1: There are two additional shafts, Shaft C and Shaft D. Blocking Shaft A or Shaft C will save all the miners if they are in Shaft A but will kill all the miners if they are in Shaft B; and blocking Shaft B or Shaft D will save all the miners if they are in Shaft B but will kill all the miners if they are in Shaft A. Suppose, first, that we are in a context accepting conditional norms in which our obligations depend on what we know. Such a context would call for a conversational background like P'' in (35).

- (35) *Evidence-dependent conversational background P'' (Miners*):*

$$\begin{aligned} P''_{AK} &= \{blA \vee blC\} \\ P''_{BK} &= \{blB \vee blD\} \\ P''_{A\bar{K}} &= P''_{B\bar{K}} = \{blN\} \end{aligned}$$

In order for (10) to be accepted in the context, all worlds in the context set where the miners are in Shaft A must determine a premise set that is compatible with our blocking Shaft A or Shaft C. So, accepting (10) would require removing world $A\bar{K}$ from the context set, since the value of P'' at $A\bar{K}$ is incompatible with the proposition $blA \vee blC$. Insofar as we don’t wish to rule out this world, this correctly predicts that (10) is rejected in our original context. (*Mutatis mutandis* for (11).)

However, suppose we shift to a context accepting conditional norms in which the act prescribed depends only on the miners’ actual locations, and not on our epistemic state, as reflected in (36).

- (36) *Evidence-independent conversational background P''' (Miners*):*

$$\begin{aligned} P'''_{AK} &= P'''_{A\bar{K}} = \{blA \vee blC\} \\ P'''_{BK} &= P'''_{B\bar{K}} = \{blB \vee blD\} \end{aligned}$$

The ‘may’ conditionals (10)–(11) are accepted in this sort of context: blA is compatible with the values of P''' at AK and $A\bar{K}$, and blB is compatible with the values of P''' at BK and $B\bar{K}$. But given that (4) is also accepted, (37) is false throughout the context set and is thus rejected.

- (4) Either the miners are in Shaft A or the miners are in Shaft B.
 (37) We mustn’t block any shaft.

For no world in the context set does the value of P''' at that world entail that we don't block any shaft.

There is no Miners Puzzle with 'must' or 'may'.

4.2 'Should'

Before turning to our examples with 'should', there is a preliminary matter that requires our attention. In §3 I argued that accepting 'Should ϕ ' doesn't require accepting that ϕ is in fact necessary. I adopted an implementation of this idea on which 'should' brackets the general indicative assumption that the worlds being talked about are in the context set. Formalizing the analysis in this way raises the question of how to model the effects of assertions of 'Should ϕ ' on the context. For any world $w \in c$, 'Should ϕ ' is true at w according to Definition 3 iff 'Must ϕ ' is true at w according to Definition 1.²⁵ So, to avoid predicting that uttering 'Should ϕ ' has the same effect as uttering 'Must ϕ ', one needs to allow that assertions of 'should' claims may distinguish among worlds outside the context set or in some relevant subset of the context set. This raises the challenge of providing an account of how exactly this process works. If asserting 'Should ϕ ' needn't involve talking about the set of live possibilities, then which worlds *are* being talked about? In which worlds is ϕ being said to be necessary?

Using a clause that lacks the standard indicative presupposition places a burden on the interpreter's task of inferring which possibilities are being talked about. Weak necessity modals, however, are neither marked as being about the actual world nor indicated as being about some other salient possibility. So, for a speaker's utterance of 'Should ϕ ' to be relevant, she must be treating the necessity claim as holding given some condition plausibly relevant to the necessity of ϕ . Recall our discussion in §3 of the conversational role of 'should'. Suppose we are considering the question of the necessity of ϕ . Uttering 'Must ϕ ' would require settling precisely which considerations apply given the circumstances. We might prefer not to restrict the future course of the conversation in this way. Nevertheless each of us takes some ways of extending the current conversation and addressing this question to be more likely or better than others. Using 'should' allows us to consider the necessity of ϕ as holding, not in the current context, as with 'must', but in a favored continuation or minimal revision of the current context, whatever that might turn out to be.²⁶

²⁵The right-to-left direction is obvious. For the left-to-right direction, if $\llbracket \text{Should } \phi \rrbracket^c(w) = 1$, then $\bigcap P_w \subseteq \llbracket \phi \rrbracket^c$; but then $\llbracket \text{Must } \phi \rrbracket^c(w) = 1$, since, by hypothesis, $w \in c$. Thanks to the editor and an anonymous referee for pressing this issue.

²⁶Compare Schueler's (2011) claim that the implicit antecedent in an implicit conditional is an an-

There may ultimately be reasons to build this feature of the use of ‘should’ into the semantics itself, rather than attempt to derive it pragmatically. That is, one might treat ‘Should ϕ ’ as true iff for all “best” (in some relevant sense of ‘best’) relevant worlds w' , ϕ follows from $P_{w'}$, where the selected worlds w' aren’t presupposed to be within the context set. I won’t take a stand on this issue here. For expository purposes I will stick with the semantics in Definition 3. Given a use of ‘Should ϕ ’, let T be the set of worlds being talked about (‘ T ’ for “Topic worlds”), the set of worlds in which ϕ is said to be deontically necessary. What is important for our purposes is that T isn’t assumed to be the context set, even after one’s assertion has been accepted. Depending on the context, T might be a subset of the context set, it might include some worlds outside the context set, or it might be disjoint from the context set. For simplicity, however, I will assume that in our examples T is a subset of the context set, representing a proposed (most likely/plausible) extension of the current context.²⁷

With these points in mind, let’s return to our examples with ‘should’ in (1)–(3), reproduced below.

- (1) We should block neither shaft.
- (2) If the miners are in Shaft A, we should block Shaft A.
- (3) If the miners are in Shaft B, we should block Shaft B.

Start with (1). Given the evidence-dependent norms encoded in P from (33), the truth of the variant with ‘must’ in (25) requires that we don’t know which shaft the miners are in.

- (25) We must block neither shaft.

Hence accepting (25) would require committing that we won’t get new evidence about the miners’ location within the interval circumscribing our deliberation (§4.1). Though we may wish to proceed as if this is indeed the case, we might prefer not to explicitly register this commitment, at least for the purposes of the conversation. Using ‘should’ provides a means of doing so. It allows us to consider a proposed extension of the current conversation in which it is accepted that we don’t learn which shaft the miners are in — as reflected in setting $T = \{A\bar{K}, B\bar{K}\}$ — and to then

tecedent that “expresses the most plausible alterations to current contextual knowledge which would make the content of the [implicit conditional] (interpreted as a consequent in the paraphrase) relevant to the context of utterance” (8; cf. KASPER 1992). See FARKAS & BRUCE 2010 on the importance of representing projected future states of the conversation in a model of discourse.

²⁷See SILK 2012b for additional discussion of these issues and various alternative implementations, including developments in a dynamic update semantics.

consider the deontic necessity of our blocking neither shaft in *that* (hypothetical) context. We can accept (1) without committing ourselves to ruling out AK and BK as live possibilities, or decisively settling that we won't get new evidence about which shaft the miners are in. Accepting (1) only requires that blN be deontically necessary at $A\bar{K}$ and $B\bar{K}$, which it is: blN follows from both $P_{A\bar{K}}$ and $P_{B\bar{K}}$. In this way, using the weak necessity modal 'should' in (1) allows us to plan under the supposition that we will remain ignorant of the miners' location, while remaining open to the possibility, however improbable, that new evidence will come in.

This treatment of the contrast between (1) and (25) predicts that if we are willing to commit that we won't learn which shaft the miners are in, then using 'must' will become more appropriate. This prediction appears to be borne out, as reflected in (17), reproduced below.

- (17) *Alice*: Well, it looks like we're never going to be able to find out which shaft the miners are in.
Bert: So what should I tell the team?
Alice: That we must block neither shaft.

We can capture how the relative felicity of unembedded 'should' and 'must' judgments depends on context.

Now turn to the 'should' conditionals in (2)–(3). I suggest that when we hear (2)–(3) as compelling, we treat the relevant set of worlds T^* as worlds in which we learn which shaft the miners are in, i.e. $\{AK, BK\}$. This would correctly predict that (2)–(3) are accepted in our original context: (2) is accepted iff blA follows from the value of P at all *inA*-worlds in T^* iff blA follows from P_{AK} , which it does. (*Mutatis mutandis* for (3).) Such a treatment of T^* in the 'should' conditionals is motivated by independent principles of interpretation. It is common to treat the consequent of a conditional as interpreted with respect to the local context set up by the antecedent — i.e., the global context hypothetically incremented with the antecedent (KARTTUNEN 1974, STALNAKER 1974, HEIM 1990).²⁸ Uttering a conditional with the antecedent 'if the miners are in Shaft A' makes salient the possibility that the miners are in Shaft A. In a context in which we are planning for various contingencies, uttering 'if the miners are in Shaft A' can also make locally salient the live possibility that we learn that the miners are in Shaft A. After all, contingency planning is plan-

²⁸Positing local ("derived," "subordinate") contexts is independently motivated by various phenomena concerning presupposition accommodation and anaphora. Though common in dynamic semantics, positing local contexts is neutral between static/pragmatic implementations (à la Stalnaker) and dynamic semantic implementations (à la Heim).

ning about what to do for the situation of *recognizing* that we are in such-and-such circumstances; it is planning for possibilities we might realize we are in. If our obligations about how to respond to the miners' plight depend on our epistemic state, a hearer can interpret utterances of (2)–(3) as appropriate by treating the topic worlds T^* as worlds in which we learn which shaft the miners are in. Principles of interpretive charity then motivate her to do so. Even if in the global discourse context it is regarded as highly unlikely that we will discover the miners' location, if it isn't settled in the conversation that we won't, then (2)–(3) can still be accepted. On these interpretations, (2)–(3) allow us to plan for the (possibly remote) possibility in which we learn that the miners are in such-and-such shaft. This reflects an important role for 'should' conditionals in contingency planning and deliberation.

Treating T^* in the 'should' conditionals in this way has several additional features. First, it makes sense of an otherwise puzzling suggestion in the literature mentioned in §2, namely, that the antecedents of the conditionals are elliptical for 'if the miners are in Shaft A (/Shaft B) and we learn it' (VON FINTEL 2012). Our account has the advantage of capturing the intuitions driving this move without having to claim that the antecedents of certain 'should' conditionals, but not those of various other conditionals, are reinterpreted.

Relatedly, second, it correctly predicts that if we explicitly settle that we won't learn which shaft the miners are in, it becomes harder to hear the 'should' conditionals as felicitous.

- (38) *Alice*: Well, it looks like we're never going to be able to find out which shaft the miners are in.
Bert: So what should I tell the team?
Alice: We should block neither shaft. ?But if they're in A, we should block A, and if they're in B, we should block B.

Alice's response to Bert is marked.

Third, we correctly predict that there should be a reading of a deontic conditional with 'we should block neither shaft' as its consequent clause, namely, if the set of relevant topic worlds is identified with the set supplied by the global discourse context. As noted above, contexts of contingency planning can make salient interpretations of (2)–(3) on which they are natural to accept. But if the primary question in the conversation is how we should plan given the most likely scenario, the relevant worlds at which the necessity of the prejacent of 'should' is evaluated will be worlds in which we *don't* learn which shaft the miners are in. In such a context, it is

the ‘should’ conditional in (39) that is more natural to accept.²⁹

(39) Even if the miners are in Shaft A, we should block neither shaft.

Finally, fourth, our treatment of the ‘should’ conditionals elucidates a previously unnoticed feature of the Miners Case: After uttering (1), it is much less natural, perhaps even infelicitous, to begin an utterance of the conditionals with ‘and’, as opposed to ‘but’.

(40) We should block neither shaft. But (/#And) if the miners are in Shaft A, we should block Shaft A; and if the miners are in Shaft B, we should block Shaft B.

Likewise for reverse sequences involving an utterance of the conditionals followed by an utterance of (1):

(41) If the miners are in Shaft A, we should block Shaft A; and if the miners are in Shaft B, we should block Shaft B. But (/#And) we should block neither shaft.

Adversative conjunctions like ‘but’, in contrast to ‘and’, are required when the second conjunct contradicts an expectation raised by the first conjunct. But if (1)–(3) are jointly consistent, as the literature has largely assumed, it is puzzling why ‘but’ should be required in (40)–(41). The above account provides a natural explanation. ‘But’ is required in order to mark that there is a shift in which live possibilities we are considering in our deliberations, a shift in which contingencies we are planning for. In (40), uttering (1) makes salient live possibilities in which we don’t learn the miners’ location, and thus leads us to expect that such possibilities will be the main topic of our deliberation and planning. Using ‘but’ before uttering the conditionals marks that this expectation is canceled. (Analogous remarks hold for (41).)

Though I have said that the ‘should’s in (1), on the one hand, and (2)–(3), on the other, are interpreted with respect to different relevant worlds, this needn’t imply that there is an equivocation in a sense that would render (1)–(4) jointly inconsistent. The shift in the conditionals results from a *local*, or sentence-internal, update. Hence, (1)–(4) can all be accepted in a constant global context. Since it is standard in assessing entailment relations to interpret the sentences with respect to a constant

²⁹A construction like ‘even if’ (‘whether or not’, etc.) is required to block the implicature or presupposition of a hypothetical conditional ‘If ψ , ϕ ’ that ‘If $\neg\psi$, $\neg\phi$ ’ is also true (GEIS & ZWICKY 1971, IATRIDOU 1991, HORN 2000). For discussion of ‘even if’ conditionals in a revisionary approach to information-sensitivity, see SILK 2014b: 696, 720.

global context (KAPLAN 1989, KADMON & LANDMAN 1993, VON FINTEL 1999), we can maintain that (1)–(4) are jointly satisfiable.³⁰ Terminology aside, what is of primary importance for present purposes is the phenomena: The ‘should’s in (1)–(3) are interpreted with respect to a constant body of norms, as reflected in a constant deontic conversational background, but are used to make necessity claims about different sets of worlds.

In sum, we correctly predict that (1)–(4) can all be accepted in our original context. By accepting the unembedded ‘should’ claim in (1), we can coordinate on a plan to block neither shaft, but without needing to settle decisively that we won’t get new evidence about the miners’ location. Remaining open to the possibility, however slight, that we might learn which shaft the miners are in, we can plan for this contingency by also accepting the ‘should’ conditionals in (2)–(3). It is the weak necessity modal ‘should’, in contrast to ‘must’ or ‘may’, that plays this complex role in conversation, deliberation, and planning.

In these sections I have offered an account of the Miners Puzzle and our additional data from §1. This account improves upon existing alternatives in empirical coverage and independent motivation (§2). But it is not the only possible account. I welcome alternative ways of systematizing and theorizing about the data with which the present account may be compared.

5 Conclusion

Discussions of information-sensitivity and the Miners Puzzle have focused on a narrow range of cases. Nearly all examples use weak necessity modals, and detailed specifications of context are rare. Many theorists have used the assumed data to motivate sweeping generalizations about the semantics of modals and conditionals. These generalizations have been premature. Expanding our data set suggests that what makes the Miners Puzzle a puzzle are features unique to weak necessity modals. Considering examples using various types of modal expressions, with targeted variations in conversational context, sheds new light on information-sensitivity, and provides a basis for a more comprehensive account of deontic modals and conditionals. The broader spectrum of examples elucidates the variety of expressive resources at our disposal for coordinating our actions and expectations in conversation, deliberation, and planning.

³⁰This isn’t to deny the importance of a dynamic notion of entailment that tracks how context gets updated in a process of reasoning. See SILK 2014b: §7 for discussion in Miners-style cases.

References

- VAN DER AUWERA, JOHAN & VLADIMIR A. PLUNGIAN. 1998. Modality's semantic map. *Linguistic Typology* 2, 79–124.
- CARIANI, FABRIZIO, STEFAN KAUFMANN, & MAGDALENA SCHWAGER. 2013. De-liberative modality under epistemic uncertainty. *Linguistics and Philosophy* 36, 225–259.
- CARR, JENNIFER. 2012. Subjective *ought*. MS, MIT.
- CHARLOW, NATE. 2013. What we know and what to do. *Synthese* 190, 2291–2323.
- COATES, JENNIFER. 1983. *The semantics of the modal auxiliaries*. London: Croom Helm.
- DOWELL, J.L. 2012. Contextualist solutions to three puzzles about practical conditionals. In RUSS SHAFER-LANDAU (Ed.), *Oxford studies in metaethics*, vol. 7, pp. 271–303. Oxford: Oxford University Press.
- FARKAS, DONKA F. & KIM B. BRUCE. 2010. On reacting to assertions and polar questions. *Journal of Semantics* 27, 81–118.
- VON FINTEL, KAI. 1999. NPI licensing, Strawson entailment, and context dependency. *Journal of Semantics* 16, 97–148.
- VON FINTEL, KAI. 2012. The best we can (expect to) get? Challenges to the classic semantics for deontic modals. Paper presented at the 2012 Central APA, Symposium on Deontic Modals.
- VON FINTEL, KAI & SABINE IATRIDOU. 2008. How to say *ought* in foreign: The composition of weak necessity modals. In JACQUELINE GUÉRON & JACQUELINE LECARME (Eds.), *Time and modality*, pp. 115–141. Springer.
- VAN FRAASSEN, BAS C. 1973. Values and the heart's command. *The Journal of Philosophy* 70, 5–19.
- FRANK, ANETTE. 1996. *Context dependence in modal constructions*. Ph.D. thesis, University of Stuttgart.
- GEIS, MICHAEL & ARNOLD ZWICKY. 1971. On invited inferences. *Linguistic Inquiry* 2, 561–566.

- HARMAN, GILBERT. 1975. Reasons. *Crítica: Revista Hispanoamericana de Filosofía* 7, 3–17.
- HARMAN, GILBERT. 1993. Stringency of rights and “ought”. *Philosophy and Phenomenological Research* 53, 181–185.
- HEIM, IRENE. 1990. On the projection problem for presuppositions. In STEVEN DAVIES (Ed.), *Pragmatics: A reader*, pp. 397–405. Oxford: Oxford University Press.
- HORN, LAURENCE R. 1972. *On the semantic properties of logical operators in English*. Ph.d. thesis, University of California, Los Angeles.
- HORN, LAURENCE R. 1989. *A natural history of negation*. Chicago: University of Chicago Press.
- HORN, LAURENCE R. 2000. From *if* to *iff*: Conditional perfection as pragmatic strengthening. *Journal of Pragmatics* 32, 289–326.
- HUDDLESTON, RODNEY & GEOFFREY K. PULLUM (Eds.). 2002. *The Cambridge grammar of the English language*. Cambr: Cambridge University Press.
- IATRIDOU, SABINE. 1991. *Topics in conditionals*. Ph.D. thesis, MIT, Cambridge.
- ISAACS, JAMES & KYLE RAWLINS. 2008. Conditional questions. *Journal of Semantics* 25, 269–319.
- JACKSON, FRANK. 1991. Decision-theoretic consequentialism and the nearest and dearest objection. *Ethics* 101, 461–482.
- KADMON, NIRIT & FRED LANDMAN. 1993. ‘Any’. *Linguistics and Philosophy* 16, 353–422.
- KAPLAN, DAVID. 1989. Demonstratives. In J. ALMOG, J. PERRY, & H. WETTSTEIN (Eds.), *Themes from Kaplan*, pp. 481–563. Oxford: Oxford University Press.
- KARTTUNEN, LAURI. 1974. Presupposition and linguistic context. *Theoretical Linguistics* 1, 181–194.
- KASPER, WALTER. 1992. Presuppositions, composition, and simple subjunctives. *Journal of Semantics* 9, 307–331.

- KATZ, GRAHAM, PAUL PORTNER, & AYNAT RUBINSTEIN. 2012. Ordering combination for modal comparison. In ANCA CHERECHES (Ed.), *Proceedings of SALT 22*, pp. 488–507. Ithaca: CLC Publications.
- KAUFMANN, STEFAN & MAGDALENA SCHWAGER. 2011. A uniform analysis of conditional imperatives. In ED CORMANY & SATOSHI ITO (Eds.), *Proceedings of Semantics and Linguistic Theory (SALT) 19*, pp. 223–238. Ithaca: CLC Publications.
- KOLODNY, NIKO & JOHN MACFARLANE. 2010. Ifs and oughts. *Journal of Philosophy* pp. 115–143.
- KRATZER, ANGELIKA. 1977. What ‘must’ and ‘can’ must and can mean. *Linguistics and Philosophy* 1, 337–355.
- KRATZER, ANGELIKA. 1978. *Semantik der Rede*. Königstein: Scriptor.
- KRATZER, ANGELIKA. 1981. The notional category of modality. In HANS-JÜRGEN EIKMEYER & HANNES RIESER (Eds.), *Words, worlds, and contexts: New approaches in word semantics*, pp. 38–74. Berlin: de Gruyter.
- KRATZER, ANGELIKA. 1991. Modality/Conditionals. In ARNIM VON STECHOW & DIETER WUNDERLICH (Eds.), *Semantics: An international handbook of contemporary research*, pp. 639–656. New York: de Gruyter.
- LASSITER, DANIEL. 2011. *Measurement and modality: The scalar basis of modal semantics*. Ph.D. thesis, New York University.
- LEWIS, DAVID. 1973. *Counterfactuals*. Cambridge: Harvard University Press.
- LEWIS, DAVID. 1975. Adverbs of quantification. In EDWARD L. KEENAN (Ed.), *Formal semantics of natural language*, pp. 3–15. Cambridge: Cambridge University Press.
- LEWIS, DAVID. 1981. Ordering semantics and premise semantics for counterfactuals. *Journal of Philosophical Logic* 10, 217–234.
- LYONS, JOHN. 1977. *Semantics*, vol. 2. Cambridge: Cambridge University Press.
- MACFARLANE, JOHN. 2014. *Assessment sensitivity: Relative truth and its applications*. Oxford: Clarendon Press.
- MCMNAMARA, PAUL. 1990. *The deontic quadecagon*. Ph.d. thesis, University of Massachusetts.

- McNAMARA, PAUL. 1996. Must I do what I ought? (Or will the least I can do do?). In MARK A. BROWN & JOSÉ CARMO (Eds.), *Deontic logic, agency, and normative systems*, pp. 154–173. Berlin: Springer-Verlag.
- MYHILL, JOHN. 1996. The development of the strong obligation system in American English. *American Speech* 71, 339–388.
- NINAN, DILIP. 2005. Two puzzles about deontic necessity. In JON GAJEWSKI, VALENTINE HACQUARD, BERNARD NICKEL, & SETH YALCIN (Eds.), *MIT working papers in linguistics: New work on modality*, vol. 51, pp. 149–178.
- PALMER, F.R. 1990. *Modality and the English modals*. New York: Longman, 2nd edn.
- PALMER, F.R. 2001. *Mood and modality*. Cambridge: Cambridge University Press, 2nd edn.
- PARFIT, DEREK. 1988. What we together do. MS, Oxford University.
- PORTNER, PAUL. 2009. *Modality*. Oxford: Oxford University Press.
- RUBINSTEIN, AYNAT. 2012. *Roots of modality*. Ph.D. thesis, University of Massachusetts Amherst.
- SCHUELER, DAVID. 2011. *The syntax and semantics of implicit conditionals: Filling in the antecedent*. Ph.d. thesis, University of California, Los Angeles.
- SILK, ALEX. 2012a. Modality, weights, and inconsistent premise sets. In ANCA CHERECHES (Ed.), *Proceedings of SALT 22*, pp. 43–64. Ithaca: CLC Publications.
- SILK, ALEX. 2012b. Weak and strong necessity. URL <http://goo.gl/xY9r0N>. MS, University of Birmingham.
- SILK, ALEX. 2013a. Modality, weights, and inconsistent premise sets. MS, University of Birmingham.
- SILK, ALEX. 2013b. *What normative terms mean and why it matters for ethical theory*. Ph.D. thesis, University of Michigan, Ann Arbor.
- SILK, ALEX. 2014a. Discourse contextualism: A framework for contextualist semantics and pragmatics. URL <http://goo.gl/aIDDmb>. Monograph, University of Birmingham.

- SILK, ALEX. 2014b. Evidence sensitivity in weak necessity deontic modals. *Journal of Philosophical Logic* 43, 691–723.
- SILK, ALEX. 2014c. How to be an ethical expressivist. *Philosophy and Phenomenological Research* .
- SILK, ALEX. 2015. What normative terms mean and why it matters for ethical theory. In MARK TIMMONS (Ed.), *Oxford studies in normative ethics*, vol. 5. Oxford: Oxford University Press.
- SLOMAN, AARON. 1970. ‘Ought’ and ‘better’. *Mind* 79, 385–394.
- STALNAKER, ROBERT. 1974. Pragmatic presuppositions. In STALNAKER (1999), pp. 47–62.
- STALNAKER, ROBERT. 1978. Assertion. In STALNAKER (1999), pp. 78–95.
- STALNAKER, ROBERT. 1999. *Context and content: Essays on intentionality in speech and thought*. Oxford: Oxford University Press.
- STARR, WILLIAM. 2010. *Conditionals, meaning, and mood*. Ph.d. thesis, Rutgers University.
- VON STECHOW, ARNIM, SVETA KRASIKOVA, & DORIS PENKA. 2006. Anankastic conditionals again. In TORGRIM SOLSTAD, ATLE GRØNN, & DAG HAUG (Eds.), *Festschrift for Kjell Johann Sæbø*, pp. 151–171. Oslo: Unipub.
- SWANSON, ERIC. 2008. Modality in language. *Philosophy Compass* 3, 1193–1207.
- SWANSON, ERIC. 2011. On the treatment of incomparability in ordering semantics and premise semantics. *Journal of Philosophical Logic* 40, 693–713.
- SWANSON, ERIC. 2012. Imperative force in the English modal system. Handout, Philosophy-Linguistics Workshop, University of Michigan.
- SWEETSER, E. 1990. *From etymology to pragmatics*. Cambridge: Cambridge University Press.
- VELTMAN, FRANK. 1976. Prejudices, presuppositions, and the theory of conditionals. In JEROEN GROENENDIJK & MARTIN STOKHOF (Eds.), *Amsterdam papers in formal grammar*, vol. 1, pp. 248–281. Central Interfaculteit, University of Amsterdam.

- WERNER, TOM. 2003. *Deducing the future and distinguishing the past: Temporal interpretation in modal sentences in English*. Ph.d. thesis, Rutgers University.
- WERTHEIMER, ROGER. 1972. *The significance of sense: Meaning, modality, and morality*. Ithaca: Cornell University Press.
- WILLIAMS, BERNARD. 1981. Practical necessity. In *Moral luck*, pp. 124–131. Cambridge: Cambridge University Press.
- WOISETSCHLAEGER, ERICH FRIEDRICH. 1977. *A semantic theory of the English auxiliary system*. Ph.D. thesis, MIT.