Deontic Modals with Complex Acts

Andrei Nasta[1,2](✉)

[1] University of Pittsburgh, Pittsburgh, PA, USA
[2] University of East Anglia, Norwich, England
and.nasta@gmail.com

Abstract. The paper identifies a key pragmatic principle that is responsible for the information-sensitivity of deontic modals. Information-sensitivity has been extensively discussed in the recent linguistic and philosophical literature, in connection with a decision problem known as the Miners' puzzle (Kolodny and MacFarlane 2010). I argue that the so-called Ellsberg paradox (Ellsberg 1961) is a more general source of information-sensitivity. Then I outline a unified pragmatic solution to both puzzles on the basis of a well-known decision procedure (MiniMax).

Keywords: Deontic modals · Information-sensitivity · Pragmatics · Acts · Outcomes · Decision-theory · Ellsberg paradox · Miners' puzzle · Minimax

What is the ideal, most desirable choice in a decision situation depends on our epistemic states. Call this the *information-sensitivity* of desirability and of similar normative concepts.

Information-sensitivity is a special form of context-sensitivity. Context-sensitivity makes the truth- or assertability-conditions of utterances dependent on features of the world. Information-sensitivity, in contrast, makes these conditions dependent on the features of the world *as known by an agent*. Often the agent does not have access to all the facts relevant for her decision, but has only limited knowledge. This knowledge can—and often does—serve the agent well in deciding what she should do. When the context makes available less than all the facts pertinent to a decision, the agent, from her limited informational standpoint, is bound to have subjective normative commitments rather than objective ones.

In this paper I am concerned with the subjective normative commitments expressed by modal vocabulary in specific contexts. I argue for an extended form of information-sensitivity of deontic modals such as *ought to* and *should* by taking a closer look at the role of certain complex actions and outcomes. I show that the information-sensitivity can be interpreted as adherence to a specific norm of decision in two well-known decision problems. While my aim is not to provide a logic of decision or a compositional semantics for deontic modals, my remarks contribute to a better understanding of the pragmatics of deontic modality, leaving the standard modal semantics untouched.

H. Christiansen et al. (Eds.): CONTEXT 2015, LNAI 9405, pp. 373–382, 2015.
DOI: 10.1007/978-3-319-25591-0_27

1 The Miners' Puzzle

A good starting point for my observation is the Miners' Puzzle (Kolodny and MacFarlane 2010). In the miners' setting (M) we have to decide for the best course of action regarding 10 miners who, as far as we know, can be in either one of two shafts, A or B. The two shafts are about to be flooded and we do not have enough sandbags to prevent this from happening. We have sandbags to block only one of the two shafts. If we block neither shaft, we save nine miners and lose one. If we block the right shaft, we lose no miner. But if we block the wrong shaft, we lose all the miners. In this context, the following sentences sound true.

(1) We ought to block neither shaft.
(2) If the miners are in shaft A (B), we ought to block shaft A (B).

To make sense of these sentences in one's modal semantics, it has been argued, one needs information-sensitive deontic modality and, more notably, the violation of a version of modus ponens. I will assume that a baseline semantics for deontic modals such as Kolodny and MacFarlane (2010) or Kratzer (2012) is on the right line, and my observation should be taken as a *pragmatic* complement to such a baseline semantics. However, I contend that information-sensitivity further affects the reasoning with deontic modals where *complex* actions—rather than simple actions, which have been the focus of discussion so far—are taken into consideration.

1.1 Complex Acts and Outcomes

In context (M), a complex action is a disjunction or conjunction of basic acts. Since the basic acts are block A (B_a), block B (B_b), and block neither shaft (B_n), complex acts are just compounds like $B_a \vee B_b$ and $B_a \wedge B_n$. Likewise for complex outcomes. What are then the predictions of the standard account for complex acts and outcomes?

With an enlarged set of actions, the desirability of the original acts change. The desirability of an act is measured by checking out whether the act chosen warrants that the miners' lives are saved in the greatest possible proportion. Interpreting S_0, S_1 etc. as the possible outcomes: 0 miners are saved, 1 miner is saved etc. we get the following entailments: $B_a \supset S_0 \vee S_{10}$ and $B_n \supset S_9$. The entailment of B_b is the same as that of B_a, since blocking shaft A (B) guarantees only a disjunctive outcome to the effect that we either lose all the miners or save all.[1] (The consequent-proposition—representing the outcome—serves to evaluate the deontic worth of the antecedent-proposition—representing the act.)

[1] In the standard modal semantics (Kratzer 2012), the consequent of such an entailment is a proposition in the ordering source, while the antecedent is the proposition in the scope of the deontic modals. Roughly, the more consequent-propositions in the ordering source follow from the antecedent-proposition, the more deontically valuable the antecedent-proposition is.

These entailments represent the original scenario, whereby B_n is preferred to either of the other two acts, because it *guarantees* S_9. The complex outcome $S_0 \vee S_{10}$, which only guarantees that 0 lives are saved, is very unappealing in comparison to the simple outcome S_9. This is indeed an intuitive result.

My question now is what happens when we consider for comparison complex acts rather than simple ones. The answer is not always straightforward. Consider, for instance, the acts $B_a \vee B_b$, $B_a \vee B_n$, and $B_b \vee B_n$. Which of these acts is best? The relevant entailments are $B_a \vee B_b \supset S_0 \vee S_{10}$ and $B_a \vee B_n \supset S_0 \vee S_9 \vee S_{10}$.[2] (The disjunction $B_b \vee B_n$ has the same entailment as $B_a \vee B_n$.) Deciding on a preference is not simple in the absence of additional theoretical assumptions, and so I will postpone discussion of this case.

However, it is much easier to decide on conjunctive acts (assuming they are available). Consider the corresponding entailment of the act of blocking both shafts, namely $B_a \wedge B_b \supset S_{10}$. Choosing to block both shafts is the best possible decision that we can imagine in setting (M), because it maximizes the lives saved (thereby minimizing the lives lost). In the miners' setting, blocking both shafts has no competitors. The other conjunctive acts, $B_n \wedge B_a$ and $B_n \wedge B_b$ are not even defined, since what would be the point of choosing e.g. to block neither shaft and to block shaft A? The act is contradictory, so useless from a deliberation standpoint.[3]

The undefinedness of certain complex acts points to a limitation of the Miners' setting. For instance, our evidence should not be restricted to cases that make most complex acts impossible. For a more comprehensive account of deontic modality we need a broader range of complex acts and outcomes. These are genuine factors in deliberation and decision making, and thus a general semantics of deontic modality should take them into account. A step towards a more general account will be taken in the following section.

2 Ambiguous Complex Acts

Daniel Ellsberg (1961) introduced a decision-theoretic setting that poses problems for some cherished principles of expected utility theory and probability theory. An Ellsberg (E) setting is the following.[4]

> **(E)** Consider a sack containing ten coloured balls. For the present decision, the relevant colours are red, blue, and yellow. The distribution of

[2] In decision-theoretic terms, complex acts and outcomes are obtained by putting together the information in the estimated desirability matrix of a decision problem. More precisely, we could gather the information on a specific row of a matrix (which corresponds to an act), or by considering for deliberation a combination of rows (acts) of the matrix (cf. Jeffrey 1965/1983). For completeness, the outcomes—seen as entailments of the acts—should be weighted with the probabilities of the relevant conditions, but we shall leave this information implicit.

[3] It is natural to stipulate that acts, viewed as propositions, should be non-empty, and thus should have at least one element (i.e., world).

[4] See also Fishburn (1986) and Halpern (2005) for further discussion.

colours is as follows. Exactly three balls are red. The remainder seven balls are blue or yellow in unknown proportion. We have to decide which ball-colour to choose, knowing that we get to keep the ball if a ball of that colour is then randomly drawn from the sack. This is desirable, since the balls are all made of massive platinum on the inside.

Setting (E), unlike (M), makes probabilistic information relevant to the decision. This context, unlike the miners' context, allows for a more meaningful deliberation about complex actions. The Ellsberg setting is more general than (M), while preserving all the properties of the (M) setting. To see this, let us reformulate the (E)-assumptions to obtain the same problem noticed in the original (M) context. To this effect, it suffices to add several assumptions which are intuitively acceptable in (E).[5]

(3) We may choose either one of R, B, or Y. (assumption)
(4) We ought to choose R (not: B or Y). (intuitive claim)
(5) We ought to choose B or Y over R or Y. (intuitive claim)
(6) If $B \geq Y$, we ought to choose B. (intuitive conditional)
(7) If $Y \geq B$, we ought to choose Y. (intuitive conditional)
(8) Either $B \geq Y$ or $Y \geq B$. (assumption)
(9) We ought to choose either B or Y. (from 6–8)

Assuming that if an act A_1 is not dispreferred ($\succeq$) to another act A_2, we ought to choose act A_1, the intermediary steps (6)-(7) seem right.[6] By disjunctive syllogism applied to (6)–(7), paired with (8), we get (9), which contradicts (4).This is the miners' problem resurfacing in (E). But the miners' problem is *independent* of the ones generated by the Ellsberg context, and the latter cannot be formulated in the (M) setting.

[5] Notation: R, B, and Y stand for red, blue, and yellow balls. $R > B$ means that there are more red balls than blue ones, or that R is more probable than B. Because the outcomes of the acts are equal, the probabilistic relation $>$ translates into a preference relation $\succ$ (but not vice-versa). So $R > B$ means that choosing red is more probable than choosing blue, but can also read as saying that the former is preferable to the latter. Finally, the disjunctions should be interpreted as inclusive, unless *either ... or*-phrases are used. E.g. opting for $R \vee B$ brings about the prize if any of the red or blue balls is then randomly drawn. (The same could be expressed in terms of conjunction, if we interpreted the letters as bets on colours, because e.g. choosing blue *or* yellow will amount to betting on blue *and* yellow.).

[6] A version of the puzzle can be stated even if we assume that the deontic necessity modal *ought to* requires a strict preference relation. To do this, we assume that one of the following should hold: $B > Y$, $Y > B$, or $Y = B$. We then formulate three conditionals having these three statements as antecedents. For instance, we'll have the new conditional: If $Y = B$, we ought to be indifferent between the three options. (The other two conditionals will be like (6)-(7), but formulated in terms of the strict relation, $>$.) It then follows by disjunctive syllogism that we ought to choose B, Y, or R, which contradicts (3).

To see one of these additional problems, it is essential to take (4) and (5) to be true.[7] This shouldn't be a problem, as most people presented with this case find them true.[8] The basic intuition is that R and $B \vee Y$ warrant clear-cut (expected) desirabilities, whereas the other acts don't. The problem appears when these preferences are coupled with the reasonable assumption that the addition of equal amounts (of probability or desirability) to each member of an inequality should not change the direction of the inequality sign. Nonetheless, if we assume *additivity*, (5) is not consistent with (4). If (4) entails that $R \succ B$, by additivity we get that $R \vee Y \succ B \vee Y$, which is exactly the opposite of (5).

Additivity is not a principle to be lightly dispensed with. Additivity is for probabilistic and decision-theoretic settings as important as modus ponens is for logical reasoning. If we think that A is more probable than B, then it is intuitive to consider $A \vee C$ more probable than $B \vee C$. Moreover, if A is more probable than B or C, it does not follow that A is more probable than both taken together. These intuitive judgements follow from additivity, and concern not only probabilities but also desirabilities.[9] Therefore, it is worth exploring the cause of non-additivity in the (E) setting, since the deontic modals in (E) are sensitive to probabilistic information. I will argue that the diagnostic of non-additivity reveals a common feature about deontic modality in both (E) and (M) settings.

The role of complex acts in our decision problem is related to their *ambiguity* status. We say, for instance, that B is ambiguous because we don't know how many blue balls there are in context (E). So we know neither the precise probability of drawing a blue ball nor its precise desirability. Y is also ambiguous. R, however, is not, because we know precisely the number of red balls, and consequently their probability and desirability. Now, by including complex acts in our decision problem, we are faced with the possibility that a non-ambiguous act can be formed from an ambiguous one, and vice-versa. This affects what we know about the probabilities involved in assessing the acts, and ultimately—as we will see—their relative desirabilities. An illustration of the shift in relative desirability is the unexpected transition from (4) to (5), or vice-versa. In possible words terminology, some worlds will be desirable/ideal when we are faced with one pool of simple acts and less appealing when we are faced with a different pool of complex acts.[10]

---

[7] Complementarity does not hold in (E) either. Complementarity requires that if $A \succ B$, then $\neg A \preceq \neg B$, where $\succ, \preceq$ etc. are preference relations. Yet $R \succ B$, but $\neg R = B \vee Y \succ \neg B = R \vee Y$, and so complementarity is violated.

[8] For experimental evidence that the pattern of reasoning is robust see references in Camerer and Weber (1992, 332ff.).

[9] In natural language semantics, additivity has been invoked as evidence for introducing a quantitative probability measure to account for (epistemic) probabilistic modals, and against the standard Kratzer-semantics of those modals (Lassiter 2010, Yalcin 2010).

[10] In Nasta (2015b) I call this property *instability* and show that it holds of preferences and deontic commitments in general.

The argument in the next section can be viewed as showing why ambiguity is problematic in (E). But the feature that generates the problem will turn out to be deeper than ambiguity, since it is present in the non-ambiguous case (M).

3 Solution: Risk Aversion

My proposal is that the information-sensitivity and the decision procedure that informs (M), also informs (E), and so that the same features produce, in both settings, problems with different principles of reasoning. In both (M) and (E) cases, we know that choosing one act ensures the overall better desirability given the outcomes in each uncertain condition. In (E) we have a fairly clear-cut preference for one of the acts (namely, choosing red). This suggests that the heart of the matter is risk aversion: certain options are preferred because their competitors involve uncertainty with respect to great losses, and produce very unappealing expected desirabilities. Certain acts are preferred even if some other acts dominate them in specific conditions,[11] essentially because dominance is offset by the undesirability of certain possible outcomes.

The short way of diagnosing our cases is to say that they implement a MiniMax strategy, a strategy that minimizes the worst losses in terms of expected value, or, in other words, maximizes the lowest expected outcome of a choice. MiniMax predicts that a unique action will be preferred in the (M) and (E) cases. It is critical in obtaining such a preference that the deliberating agent does not have the information needed to raise the desirability of the alternative actions which incur huge risks. So we can interpret the information-sensitivity of deontic modals as follows: *what we have to do depends on our knowledge, which must exclude or minimize the possibility of the most undesirable outcomes.*[12]

The problems concerning reasoning-principles appear because (M) and (E) implicitly introduce assumptions inconsistent with minimizing the worst outcomes. (M) introduces the problematic assumption through an application of modus ponens. (E) introduces the problematic assumption through an application of additivity.

Take additivity first. This principle requires that $A \succ B$ entail $A \vee C \succ B \vee C$, and vice-versa. For additivity to respect a decision procedure (e.g. MiniMax), it has to be the case that the introduction of disjunction in each member of the inequality ($\succ$), does not change their relevant decision-theoretic status needed by the decision procedure. What is this status? We may call this status the *acceptable desirability* of a proposition reflecting a decision, which, in line with MiniMax, ensures the minimal maximal loss (or minimal lowest gain) by committing to that proposition (or act). Thus, inference must preserve not simply certain probabilistic/truth values (as required by logic and probability theory), but also acceptable desirability (as required by the decision procedure). The

[11] E.g. in (M), blocking shaft A dominates the other two acts under the condition that the miners are in shaft A.

[12] This decision procedure is closely related to the minimax regret rule invoked in rational choice theory; see Levi (1980, 144ff.) for discussion and references.

acceptability status depends on the contextually relevant decision procedure, and in (E) that procedure is MiniMax. Thus, the introduction of disjunction should preserve this desirability status.

However, disjunction does not generally preserve *acceptable desirability*, and consequently the principle of additivity—which relies on disjunction introduction—is not valid relative to the MiniMax procedure. For instance, if I prefer R to B in (E), and thus find R acceptably desirable, it does not follow that $R \vee Y$ is acceptably desirable. It should now become clear where the inconsistency lies. The inconsistency is generated by the ambivalence with respect to the MiniMax procedure in reasoning about the (E) setting. On the one hand we assume MiniMax to conclude that R is desirable, and, on the other hand, we implicitly violate MiniMax to derive the desirability of $R \vee Y$. But since disjunction introduction—essential to additivity—does not preserve acceptable desirability, and MiniMax requires just that, either additivity or MiniMax should be given up.[13] It is easy to see that the common judgement of the Ellsberg cases inclines us towards keeping MiniMax. After all, very robust intuitions in (E) lead us to minimize the greater losses in desirability.

In terms of information-sensitivity, the trouble with the additivity-based reasoning is that acceptable desirability presupposes a certain knowledge, namely knowledge of the acceptable desirability of a certain proposition. But knowledge of the acceptable desirability of propositions is not closed under disjunction (unlike preservation of truth). That is, in establishing deontic claims, we cannot rely on 'unacceptably desirable' worlds, i.e., worlds in which the acts are not known to be acceptably desirable, or worlds in which, for all we know, the worst outcome occurs. So we cannot infer deontic modal claims from propositions that contain unacceptably desirable worlds as their denotations.

What trouble additivity makes for (E), modus ponens makes for (M). This time, the trouble is not that acceptable desirability is not preserved, but that modus ponens introduces a false assumption of acceptable desirability. In terms of information-sensitivity, the conditional (that is needed in the modus ponens inference) introduces the assumption that it is known that a certain proposition (characterizing a desirable outcome) has the acceptable desirability status, which means that it is known to the deliberating agent that proposition has a maximal minimal desirability.

The assumption that the deliberating agent has information that rules out the worst outcomes is false, and reasoning on its basis generates a contradiction. As before, the contradiction is generated by applying MiniMax to obtain a conclusion, while using in another part of the reasoning an inference rule (viz. modus ponens) which is incompatible with MiniMax. The inference rule thus yields a conclusion which is inconsistent with the MiniMax conclusion independently derived.

As suggested, the latter problem is present in both (E) and (M) settings. Note that the antecedents of the conditionals in the two puzzles, (E) and (M), update the modal background with epistemic information. Thus, we get the following readings.

[13] I discuss more extensively the role of disjunction in settings (M) and (E) in Nasta (2015a).

(10) a. If the miners are in shaft A, we ought to block shaft A.
b. If the miners are in shaft A, *and we know it*, we ought to block shaft A.

(11) a. If $B > Y$, we ought to choose B.
b. If $B > Y$, *and we know it*, we ought to choose B.

If the emphasised implications[14] wouldn't go through, and we would still want the conditional to be good in the context of deliberation, we need to modify the antecedents. Accordingly, even if the miners are in shaft A, we still need to block neither shaft. (And similarly for the Ellsberg case.) In other words, minimizing risk would require us, as decision makers, to take an option which is optimal with respect to our state of information, but would be sub-optimal with respect to a richer—and, by hypothesis, unavailable—body of information.

My diagnostic for the problematic reasoning is that when we are going through e.g. the miners' scenario, we have a proclivity for erroneously assuming, for the sake of local coherence, that the deliberating agent knows which shaft is open. So we simply add the assumption of relevant knowledge to our information state, though this assumption is not justified by the facts in the global context. Note that the proposed diagnostic—based on a local, pragmatically triggered error—does not amount to an error theory. First, such local pragmatic implications are useful in communication, and seldom generate problems. (No cooperative speaker would reason her way through these decision problems in the way indicated here!) More importantly, our diagnostic is consistent with the observation that competent speakers and reasoners become quickly aware of the tension between the assumptions in the local and global contexts. However, acknowledging this local coherence-based reasoning makes it less surprising that we can momentarily fall pray to such inconsistency.

To sum up, both (M) and (E) in their conditional formulations give rise to problems when modus ponens is applied without regard to the MiniMax decision procedure. In addition, the application of the additivity principle in (E) determines a further violation of the MiniMax decision procedure. The violation of MiniMax has been traced back to the introduction of disjunction (in the case of additivity), and to a coherence-based implication of the conditional (in the case of modus ponens). Since in the present context it is plausible that MiniMax is being used in deriving the preferred act (or the corresponding deontic modal proposition), violations of MiniMax in other parts of the reasoning generate inconsistency. Thus, taking MiniMax as fundamental, we have an explanation of the puzzling inconsistencies obtained in the (M) and (E) settings.

The information-sensitivity of deontic modals can be characterized by their contextual sensitivity to the decision procedure, which in turn requires knowledge. As we have seen, the truth-value of a deontic modal claim depends on the decision procedure used to establish that claim, which amounts to saying that an *ought*-claim is sensitive to what we know to be the best solution (according

[14] See von Fintel (2012, pp.28–29) for relevant evidence of this type. This evidence suggests that conditionals admit of implicit restrictors which are sensitive to local pragmatic implications. Such local pragmatic (coherence-based) implication exists in several other linguistic domains (cf. Simons 2014).

to the decision procedure). But what we know and what we *assume* to know in making inferences can come apart, since certain inferences may introduce epistemically and deontically unwarranted assumptions, as they do in our cases.

4 Concluding Remarks

In light of the previous discussion, I agree with Kolodny and MacFarlane (2010, pp.130,136) that what we ought to do depends on our knowledge. I also agree with them that modus ponens does not satisfy this desideratum. However, I have a more specific take on the sort of knowledge that matters and how the illicit assumption of relevant knowledge comes about in the case of conditionals. In contrast to previous accounts of the miners' puzzle—e.g. Kolodny and MacFarlane (2010), Cariani et al. (2013), Charlow (2013)[15]—my approach is more explicit on how to apply the decision-theoretic principles to contexts involving probabilistic information.

It is well beyond the scope of this paper to offer a complete recipe for picking out contexts where MiniMax is successfully applicable. It is nonetheless worth asking, How general is the strategy proposed here? My contention is that the strategy is more general than we might have guessed by looking at an isolated decision setting. At the very least, it applies to cases where some outcomes are certain (as in M) and to cases where the outcomes are not certain (as in E), irrespective of whether the actions evaluated are simple or complex, and (arguably) irrespective of their ambiguity status. However, I cannot make claims about the precise bounds of the cases for which the strategy would work. As I suggested, decision procedures other than MiniMax are relevant in other different settings. We cannot build preferences based on a unique decision procedure. Moreover, even keeping fixed the broad context of e.g. the (M) or (E) case, it's not clear that MiniMax should apply to whatever decision problem we may come up with, and indeed it might well be indeterminate whether any decision procedure applies at all.[16]

To conclude, my proposal gives pride of place to a decision norm in the pragmatics of deontic modality in a particular type of context, as provided by the (M) and (E) cases. In this type of context, I interpreted the information-sensitivity

[15] Though see Carr (2012) and Lassiter (2011), who directly approach the miners' puzzle, and Goble (1996) whose deontic logic account may constitute a good framework for dealing with both (M) and (E).

[16] On the one hand, it is easy to check that MiniMax gets a good prediction for the conjunctive act of blocking both shafts in (M), and for the act of choosing a red ball in (E). On the other hand, it is much more difficult to come up with a sharp comparison between the disjunctive acts $B_a \vee B_b$, $B_a \vee B_n$, and $B_b \vee B_n$. The latter two acts guarantee that either no miner will be saved or nine miners will be saved or all of them will be saved ($S_0 \vee S_9 \vee S_{10}$), whilst the former act guarantees that either zero or ten miners will be saved ($S_0 \vee S_{10}$). A first problem is that in order to estimate the desirabilities of the acts we have to come up with probabilities for the disjuncts. And even after doing that, it is not clear that there will be only one obvious way of choosing between the estimated desirabilities thus obtained.

of deontic modality as sensitivity to a MiniMax decision strategy. Deontic modal talk in such contexts is sensitive to what might be the worst outcomes. If our information state leaves the occurrence of the worst outcome open, the act leading to that outcome is dispreferred. This holds true for both simple and complex acts, and my introduction of (E) was in part motivated by the observation that (M) somewhat obscures this fact. Nonetheless, my primary motivation has been to provide a common diagnostic for the problems raised by (M) and (E). In virtue of capturing this unifying feature, my analysis offers guidelines for a more comprehensive pragmatics of deontic modals under uncertainty.

References

Camerer, C., Weber, M.: Recent developments in modeling preferences: uncertainty and ambiguity. J. Risk Uncertainty **5**, 325–370 (1992). (cit. on p. 5)

Cariani, Fabrizio, Kaufmann, M., Kaufmann, S.: Delib- erative modality under epistemic uncertainty. Linguist. Philos. **36**, 225–259 (2013). doi:10.1007/s10988-013-9134-4. (cit. on p. 9)

Carr, J.: Subjective Ought. ms. MIT (2012). (cit. on p. 9)

Charlow, Nate: What we know and what to do. Synthese **190**, 2291–2323 (2013). doi:10.1007/s11229-011-9974-9. (cit. on p. 9)

Ellsberg, D.: Risk, ambiguity, and the Savage axioms. The Quarterly Journal of Economics, pp. 643–669 (cit. on pp. 1, 3) (1961)

Fishburn, P.C.: The axioms of subjective probability. Stat. Sci. **1**(3), 335–358 (1986). (cit. on p. 3)

Goble, L.: Utilitarian deontic logic. Philos. Stud. **82**(3), 257–317 (1996). (cit. on p. 9)

Halpern, J.Y.: Reasoning About Uncertainty. MIT Press, Cambridge (2005). (cit. on p. 3)

Jeffrey, R.C.: The logic of decision, 2nd edn. University of Chicago Press, Chicago (1965/1983) (cit. on p. 3)

Kolodny, N., John, M.: Ifs and oughts. J. Philos. **107**(3), 115–143 (2010). (cit. on pp. 1, 2, 9)

Kratzer, A.: Modals and Conditionals. Oxford University Press, Oxford (2012). doi:10.1093/acprof:oso/9780199234684.001.0001. (cit. on p. 2)

Lassiter, D.: Gradable epistemic modals, probability, and scale structure. In: Li, N., Lutz, D. (eds.) Semantics and Linguistic Theory (SALT), vol. 20, pp. 197–215. CLC Publications, Ithaca (2010). (cit. on p. 5)

Lassiter, D.: Measurement and Modality: the Scalar Basis of Modal Semantics. Ph.D. thesis, New York University (2011) http://semanticsarchive.net/Archive/WMzOWU2O/ (cit. on p. 9)

Levi, I.: The Enterprise of Knowledge. MIT Press, Cambridge (1980). (cit. on p. 6)

Nasta, A.: Disjunctive deontic modals. ms. Pittsburgh/East Anglia (2015a) (cit. on p. 7)

Nasta, A.: Unstable preferences, unstable deontic modals. ms. Pittsburgh/East Anglia (2015b) (cit. on p. 5)

Simons, Mandy: Local pragmatics and structured contents. Philos. Stud. **168**, 21–33 (2014). doi:10.1007/s11098-013-0138-2. (cit. on p. 8)

von Fintel, K.: The best we can (expect to) get? Challenges to the classic semantics for deontic modals. In: 85th Annual Meeting of the American Philosophical Association, Chicago (2012). http://web.mitedu/fintel/fintel-2012-apa-ought.pdf (cit. on p. 8)

Yalcin, S.: Proabability operators. Philos. Compass **5**(11), 916–937 (2010). doi:10.1111/j.1747-9991.2010.00360.x. (cit. on p. 5)