

Epistemic and Deontic *Should*.

Fabrizio Cariani
Northwestern University

[This is a postprint of a paper published in *Thought*, 2(1), pp. 73-84, 2013.
Submitted: 1/22/13. Accepted: 6/25/13. First Published Online: 8/13/13]

I

Sentences involving *should* and its near-synonym *ought* can have epistemic as well as deontic interpretations. Contrast:

- (1) There should be beer in the fridge.
- (2) You should bet on Barcelona.

Sentence (1) expresses a relation between the prejacent proposition (the proposition that there is beer in the fridge) and a body of evidence. Similarly, (2) expresses a relation between a goal (presumably, winning money) and its own prejacent proposition (the proposition that you bet on Barcelona). Furthermore, these relations appear to be probabilistic, thus giving some support to recent attempts to formulate probabilistic theories of the meaning of *should* and *ought* (Finlay 2010, Lassiter 2011, Yalcin 2012). When we try to specify these truth-conditions, however, a puzzle arises.

The puzzle is that the relevant relations are structurally quite different. In the epistemic case, our judgments track the probability of the prejacent given the evidence. For example, (1) may be paraphrased as:

- (3) Given the contextually salient evidence, it is probable that there is beer in the fridge.

By contrast, in the deontic case, what seems relevant is the probability of contextually given goals given the prejacent.

- (4) Given that you bet on Barcelona, it is probable that contextually salient goal(s) are satisfied.

As an analysis of (2), (4) is rather incomplete and unsatisfying. But it is a decent first stab at systematizing the thought that our inclination to accept (2) varies with the probability of a Barcelona win. It would be reasonable for a probabilistic account of deontic *should* to be a refinement of (4). In any case, (1) clearly depends on the probability of its prejacent in a way that (2) doesn't.

A similar asymmetry arises with the following pair of inferences:

- (5)
 - a. If Larry shopped today, there should be beer in the fridge.
 - b. Larry certainly shopped today.
 - c. So, it's probable that there is beer in the fridge.
- (6)
 - a. If Larry shopped today, Mary shouldn't shop.
 - b. Larry certainly shopped today.
 - c. So, it's probable that Mary didn't shop.

I doubt that either inference is valid, but (5) is pretty strong, while (6) is clearly bad. This difference in strength needs an explanation: a plausible one is that epistemic, but not deontic *should* is sensitive to the probability of the prejacent.

If these considerations are correct, the probabilistic relations most naturally associated with epistemic and deontic readings are distinct. Since it is otherwise desirable to keep the meanings of epistemic and deontic *should* as unified as possible, this asymmetry spells trouble for probabilistic accounts of *should* (and *ought*). Specifically, the puzzle appears to set up the following theoretical choice.

- (i) to keep the treatment of deontic *should* in (4) and assimilate the epistemic readings to it.
- (ii) to keep the treatment of epistemic *should* in (3) and assimilate the deontic readings to it.
- (iii) to refrain from allowing probabilistic structure in modal semantics.
- (iv) to give a bifurcated account—i.e. offer a different semantics for epistemic and deontic readings.
- (v) to account for both epistemic and deontic uses in terms of expected values (Lassiter 2011).

I will argue that there is an attractive picture of the relationship between epistemic and deontic *should* that doesn't endorse any of these options. Before advancing it, I will raise some serious objections against (i) and (ii).

I am not going to argue against (iii), even though it is the dominant theoretical picture (for example, the semantic theory associated with the work of Kratzer 1981, 2012 does not deploy probabilistic tools). There is enough intuitive evidence, even just from our judgments about (1) and (2), to at least entertain the idea that *should*-sentences involve probabilistic information and to justify exploring theories that assign some explicit role to such information.

Once my proposal is spelled out, it will be clear that (iv) isn't the correct response to this particular puzzle (although, for all I say here, a bifurcated account might be motivated in different ways). It will also be clear that there are advantages of my proposal over accounts in the style of (v). I won't dwell too much on

these, since a complete comparative assessment must be left for separate work. The central aim of this note is to establish an important possibility result, not to argue the superiority of my account over all rivals.

II

Against (i): according to (i), the truth-conditions of sentences involving unembedded epistemic *should* track the conditional probability $Pr(\textit{evidence}|\textit{prejacent})$. In our example, the truth conditions of (1) would roughly be given by:

(1') Given that there is beer in the fridge, the contextually salient evidence is probable.

This must be mistaken: if the unconditional probability of the evidence is 1, then for any prejacent, $Pr(\textit{evidence}|\textit{prejacent})$ must also be 1 (as long as the prejacent has non-zero probability). Yet, even when the unconditional probability of the evidence is 1, not all *should*-sentences are acceptable.

Against (ii): on the converse picture, deontic *should* tracks $Pr(\textit{prejacent} | \textit{goal})$. For example, one might say that *Should* φ is true iff the proposition expressed by φ is sufficiently probable (“sufficiently probable” could be spelled out in terms of the conditional probability of the alternatives given the goal). In our examples, the truth conditions of (2) would roughly be given by:

(2') Given that the contextually salient goals are satisfied, it is sufficiently probable that you bet on Barcelona.

This rendition is problematic when the probability of one’s betting on Barcelona is independently high. Consider Sam, a superstitious Barcelona fan with a passion for gambling: he only bets against Barcelona if he sees an intricate and improbable cloud pattern in the sky. Otherwise, he bets for them to win. Despite this, his goal in betting is unequivocally to win money. Suppose that the cloud pattern is probabilistically independent of the result of the game. In such a case, $Pr(\textit{prejacent} | \textit{goal})$ is extremely high: given that the goal (winning money) obtains, it is extremely likely that Sam bets on Barcelona. The reason for that, however, is just that it is independently likely that Sam bets on Barcelona and this event is independent of a Barcelona victory. Imagine, further, that this is a game in which Barcelona is slightly more likely to lose than they are to win, so that Sam should not bet on Barcelona. Option (ii) implies, contrary to intuition, that (2) is true in such a setting.

The general moral of this kind of case is that $Pr(\textit{prejacent} | \textit{goal})$ may be high for a variety of reasons—even ones that are irrelevant to the truth of the deontic

claim. In an extreme example, when the prejacent and the goal are probabilistically independent, option (ii) makes *Should* φ depend directly on $Pr(\phi)$, which seems mistaken.

Finlay (2010) defends an approach along the lines of (ii). His way around this problem is to adopt a further assumption about the structure of the probability function(s) involved in deontic readings:

Symmetry of Choice: given a set of contextually salient alternatives A_1, \dots, A_n , the probability of each A_i is identical.

In Finlay’s words: “the base should involve an assignation of equal initial probability to the choice of every relevant potential means” (p. 80). Under Symmetry of Choice (as long as the prejacent proposition is one of the A_i ’s), there is no difference between:

- comparing alternatives according to $Pr(goal \mid prejacent)$ and
- comparing them according to $Pr(prejacent \mid goal)$.

That is, Symmetry of Choice collapses the distinction between the two ways of setting up the truth-conditions of deontic *should*.¹

I reject Symmetry of Choice for two reasons. First, it is plausible to assume that the contextually supplied probability function might be either some salient credence (e.g. the agent’s, or the speaker’s) or an evidential probability function (by which I mean: a probability function that is objectively determined by a body of evidence, in the sense developed by Williamson 2000, pp. 209-37).² Given the evidence in the superstitious supporter case, an evidential probability function ought to imply that a bet on Barcelona is more probable than the alternatives. Furthermore, it is permissible for the conversational participants to have subjective credences that also imply this verdict (i.e. that a bet on Barcelona is more probable). In fact, I could stipulate it to be part of the case that the subjective credences match the evidential probability on the relevant propositions. Going

¹More precisely, let A and B be variables ranging over alternatives. Let G be a goal. Say that $A >_1 B$ iff $Pr(A \mid G) > Pr(B \mid G)$ and that $A >_2 B$ iff $Pr(G \mid A) > Pr(G \mid B)$. Under Symmetry of Choice, these orderings must be coextensive. *Proof*,

$$\begin{aligned} Pr(A \mid G) > Pr(B \mid G) & \text{ iff } Pr(A \& G) > Pr(B \& G) \\ & \text{ iff } Pr(A \& G) / Pr(A) > Pr(B \& G) / Pr(B) \\ & \text{ iff } Pr(G \mid A) > Pr(G \mid B) \end{aligned}$$

The second equivalence holds here because A and B are both alternatives, and hence by Symmetry of Choice $Pr(A) = Pr(B)$. The reasoning is the same if the sign is ‘=’.

²To clarify, I have not claimed that it *has to be* of one of these kinds: for all I say, deontic claims might be sometimes evaluated relative to more objective probability functions. My point is just that these seem to be *permitted* choices.

beyond the superstitious supporter case, it seems clear that Symmetry of Choice won't be satisfied by either subjective or evidential probabilities in a wide variety of cases. I take this to be a strike against Symmetry of Choice.

A second problem arises by the interaction of Symmetry of Choice and deontic conditionals:

(7) If it's snowing outside, you should take the train.

The evaluation of deontic conditionals on a probabilistic account is a delicate matter (Yalcin 2012, Cariani, ms.), but for basic cases such as (7), the idea is pretty simple: (7) is true in a context c iff:

(8) You should take the train.

is true relative to a context c' whose probabilistic coordinate is updated with the information that it is snowing outside.

The problem for Symmetry of Choice is that there is no guarantee that a set of alternatives equiprobable relative to an initial background will remain equiprobable after that background is updated. Suppose, for example, that according to the initial probability function, $Pr(driving) = Pr(taking\ the\ train)$. This equality may hold unconditionally, but it needn't hold after we update with the antecedent of (7). Perhaps a large proportion of snow-worlds make it impossible for you to drive. So, even assuming that it holds for unembedded sentences, Symmetry of Choice might fail in certain embeddings. This is both troubling in general and insufficient to explain our asymmetry when it shows up in embeddings (as in the inferences (5)-(6)).

Summing up, Symmetry of Choice seems to be grounded on a plausible intuition: that how we rank a deontic alternative A depends, in part, on $Pr(goal | A)$ and does not depend on the prior probability of A . But there is a big gap between claiming that the prior probability of the alternatives is irrelevant and stipulating that all alternatives ought to have the same probability. My proposal aims to save the former fact, without the latter, dubious commitment. Meanwhile, I think we should reject Symmetry of Choice, and with it the ambition to rescue option (ii).

III

To solve the puzzle I started with, I propose a framework that allows a novel kind of interaction between deontic modals and probabilities.³ Before explaining my account, consider first a classical, non-probabilistic theory:

³The framework generalizes the non-probabilistic approach in Cariani, Kaufmann and Kaufmann (forthcoming).

Should φ is true iff φ is true at all the worlds that are maximally ranked by a contextually supplied ordering.⁴

The ordering needn't be directly supplied by context; it can (as in the theory of Kratzer 1981, 2012) be built out of premises, which are in turn supplied by context (for this reason, this is called *premise semantics*). On a classical picture, premises are sets of worlds. In the deontic case, premises are understood as contextually salient goals, desires, commitments, values, etc. (following Portner 2009, call them *priorities*). For example, the goal of winning a bet may be represented by the priority:

(9) $\{w \mid \text{you win the bet in } w\}$

If (9) is the only priority, the classical theory says you should bet on Barcelona iff you bet on Barcelona at all the worlds at which you win your bet. If there are multiple, possibly conflicting, priorities, the classical recipe for ordering worlds is:

$w > v$ iff w satisfies every priority that v satisfies and some priorities that v does not satisfy.

I propose two structural changes to the classical approach. Instead of ordering worlds, I order alternatives that are supplied by context—e.g. *bet on Barcelona*, *bet on Arsenal*, *do not bet*.

This kind of contrastive theory is familiar in the semantics for deontic modals.⁵ For such theories, sentences like (2) are evaluated against a set of contextually supplied alternatives. In the deontic case, alternatives are understood as possible choices available to the salient agent(s) and modeled as sets of worlds. This move immediately invites my second change: representing priorities as sets of alternatives (rather than sets of worlds). The following are examples of priorities in this sense:

(10) $\{A \mid \text{given } A, \text{ it is probable that you win the bet}\}$

(11) $\{A \mid A \text{ guarantees that you will win the bet}\}$

(12) $\{A \mid A \text{ is compatible with your winning the bet}\}$

Notice that it is very difficult to draw the distinctions introduced by (10), (11) and (12) in a framework that treats priorities as sets of worlds (after all, a world v is compatible with your winning the bet iff it guarantees that you win the bet).

⁴Here and everywhere below, I make the limit assumption—i.e., the claim that in a partial order, every linearly ordered sub-chain terminates in a maximal element.

⁵See the analyses in Sloman (1970), Jackson (1985), Jackson and Pargetter (1986), Horty (2001), Snedegar (2012), Cariani (2013).

To distinguish them from the classical priorities, I call these ‘elevated priorities’ (although if the context is unambiguous, I occasionally revert to ‘priorities’).

Elevated priorities generate an ordering of alternatives similar to the classical ordering of worlds. Let A, B be alternatives. Then:

$A > B$ iff A satisfies every priority that B satisfies and some priorities that B does not satisfy.

There is a twist however: when we evaluate whether an alternative, say *Betting on Barcelona*, satisfies a priority, we don’t consider every possibility in which you bet on Barcelona, but rather only those possibilities that are compatible with the salient information (e.g. only those possibilities in which Barcelona is an excellent, fast-paced team, with superb fitness, etc.). A better formulation of my ordering formation recipe would then be:

$A > B$ (relative to a background set of worlds I) iff $A \cap I$ satisfies every priority that $B \cap I$ satisfies and some priorities that $B \cap I$ does not satisfy.

Informally, only the contextually salient part of A affects its position in the ordering (in a given context).⁶ This means that the very same priority π might apply to A relative to a background I , but not relative to another background I' .

We might use this ordering to replicate the classical approach and give a quantificational semantics (for more detail, see section V):

Should φ is true iff φ is true at every world that belongs to one of the maximally ranked alternatives.

Alternatively, and without much complication, we can give a non-quantificational semantics. The non-classical account of Cariani (2013) only requires a set of alternatives, a ranking of alternatives and a ‘benchmark’ that captures the options that can undermine the truth of a deontic *should*-sentence. In this note, I only develop the classical variant.

There are significant independent advantages to treating priorities as sets of alternatives. To start, the resulting approach is attractively modular and flexible: probabilistic information can affect deontic orderings, but it does not have to; a deontic ordering could also be generated by purely qualitative priorities. Another advantage is that the resulting semantic theory doesn’t encode a specific probabilistic relationship between alternatives and priorities, which allows us to implement various deontic principles without committing to a specific substantive theory of what an agent should do (for discussion of this point, see Carr *ms.*,

⁶I assume, as in the classical semantics, that this background could be set by some salient information state, but that it can also be set by in some more objective way.

and Cariani *ms.*). In this respect, my account is remarkably more neutral than theories based on expected value (it does not involve a numerical assignment of value to individual worlds, and it allows the modeling of theories such as the dominance-based theories of Horty 2001). Whether this is a genuine advantage is a matter for a separate discussion, but this kind of neutrality *has* been a hallmark feature of formal semantic theories of modality and it seems to be a feature worth preserving. Finally, the framework gives a clear, and novel, implementation of the property that Kolodny and MacFarlane 2010 call *serious information dependence*: that is to say, the ordering of alternatives may change as our information changes even if the priorities remain constant (see section V, and my Cariani *ms.*).

IV

For our purposes, the crucial advantage of this account is that it generalizes to epistemic *should* in a way that allows us to solve our puzzle. We must first establish that, like deontic *should*, epistemic *should* is sensitive to the relevant alternatives. To this end, I adapt an argument by Yalcin about *probably* (2010, p.931; the argument is based on experimental work by Windschidtl and Wells, 1998). Imagine a fair raffle with 1000 tickets; suppose that Bloggs owns 420 tickets and consider:

(13) Bloggs is probably the winner of the raffle.

We tend to accept (13) if the background is that 580 people hold one ticket each; but we reject it if a single person owns all the remaining tickets. This datum strongly suggests that *Probably* φ must mean something roughly like ‘ φ is more likely than the alternatives’ or at any rate that it expresses some kind of comparison between the prejacent of *probably* and the salient alternatives.

Yalcin’s argument applies with equal force to epistemic *should*. Consider:

(14) Bloggs should win the raffle.

Sentence (14) patterns with (13) in being sensitive to how we identify the salient alternatives. In particular, there is a true reading of (14): it arises if we contrast Bloggs with other potential winners in the scenario in which 580 people have a ticket each. If (13) supports a contrastive treatment of *probably*, (14) supports a contrastive treatment of epistemic *should*.⁷

⁷There is a passing remark in Yalcin (2012) to the effect that deontic *ought* is less sensitive to alternatives than epistemic *ought*. I have not seen Yalcin’s evidence for this claim, but I doubt that it is incompatible with the alternative-sensitivity I advance here. Suppose that Bob has a coupon to place a bet on possible winners for the raffle. Now consider the statement:

Importantly, extending Yalcin’s argument to *should* does not commit us to the view that *should* and *probably* are synonymous. An anonymous referee points out two key differences. First, suppose I am nearly hit by a car; the following sentence licenses *should* but not *probably*:

(15) That was so close! I should be / probably am* dead right now.

Second, *should* does not seem to give rise to the puzzle discussed by Yalcin (2007). Contrast:

- (16) a. Suppose Bill is sick but he should not be...
b. *Suppose Bill is sick but he probably isn’t...

The supposition in (16-b) crashes, but the one in (16-a) doesn’t. It is beyond the scope of this note to understand why this discrepancy arises, but noticing it suffices to agree that we should not exaggerate the similarities between *should* and *probably*.

As in the account of deontic *should*, the truth-conditions of epistemic *should* depend on an ordering of epistemic alternatives generated by a set of premises. The main difference is in the kinds of premises we use. The most natural option, though by no means the only one, might be:

(17) {A | A is more likely than the alternatives}

Recall that, on the present picture, when (17) is applied to an alternative *A*, it picks out only the contextually salient “portion” of *A*. So, as our information varies, we have corresponding changes in which alternatives satisfy (17). Once we interpret epistemic *should* in this way, (1) gets assigned truth-conditions roughly like those in (3).

Taking stock, the central hypothesis is that the distinction between epistemic and deontic readings of *should* can be represented as stemming from the kind of premises we use to rank alternatives. This idea should match the views of advocates of the classical premise semantics. I contribute, however, a framework that allows us to entertain the same explanation while at the same time allowing probabilistic structure to affect the truth-conditions of *should*-sentences.

V

To conclude, I offer a more precise implementation of the ideas I sketched. The reader should keep in mind that the space of options is broader than this sketch

- (i) Bob should/ought to bet on Bloggs.

Clearly, (i) seems acceptable if 580 people hold one ticket each; it seems unacceptable if a single person holds all 580 of the tickets.

suggests (this larger space is explored in Cariani *ms.*) and that many important details are intentionally omitted.

Suppose we interpret sentences relative to a point of evaluation consisting of a world w and a context c ; c consists of:

- a set of worlds i_c representing an information state
- a set of alternatives Alt_c
- a set of elevated priorities Π_c (i.e. a set of sets of alternatives).
- a probability function Pr_c

(The subscript on each coordinate keeps track of the context it is associated with, e.g. i_c is the information state of c). In this sketch, we won't worry about shifting these coordinates—although any complete theory must allow some shifting, e.g. by conditional antecedents.

Relative to c , we can order alternatives by:

$$A \geq_c B \text{ iff } \{\pi \in \Pi_c \mid (A \cap i_c) \in \pi\} \supseteq \{\pi \in \Pi_c \mid (B \cap i_c) \in \pi\}$$

Note that the ordering of alternatives only depends on c and not also on w .⁸

To implement the quantificational view, define a domain for *should* as:

$$Domain_{should}(c) = \{w \mid \exists B \in Alt_c[\sim \exists A \in Alt_c(A >_c B \ \& \ w \in B)]\}$$

Using the standard notation $\llbracket \cdot \rrbracket^{c,w}$ for the truth-value of a sentence at a point of evaluation, and $\llbracket \varphi \rrbracket^c$ for the set of worlds at which the proposition expressed by φ in c is true, we can say:

$$\llbracket should \ \varphi \rrbracket^{c,w} = True \text{ iff } \forall v \in Domain_{should}(c), v \in \llbracket \varphi \rrbracket^c$$

As mentioned, it is easy to develop a non-quantificational version of the account, though I won't do it here.

Everything that was claimed in III holds in this implementation. For example, I claimed that the current approach is information-dependent in the sense of Kolodny and MacFarlane (2010). To see this, we must first characterize information-dependence in the current setting. Given a proposition α , let $c + \alpha$ be the result of intersecting i_c with α and keeping every other coordinate of c unchanged (intuitively, this captures one way of updating c with factual information). Information-dependence then is the constraint:

$$(ID) \text{ There are } c \text{ and } \alpha, \text{ s.t. } Domain_{should}(c) \cap \alpha \not\subseteq Domain_{should}(c + \alpha)$$

⁸The ordering might be made to depend on w if, as seems plausible, we want *should* judgments to be contingent.

Informally: there are contexts in which some world v is part of the domain of *should* and also belongs to proposition α ; yet v does not belong to the domain after we update c with α . Remarkably, (ID) allows the domain to change even though the priorities remain stable ($c + \alpha$ does not modify Π_c).

An example can reveal why this constraint is satisfied in the present implementation. Suppose you go to a cheap restaurant. The house will give you whatever sauce they have today: it's either red or white sauce and the probability of each is .5. You must choose a pasta to go with the sauce. There is .5 probability that you will like spaghetti with red sauce. You will certainly like spaghetti with white sauce and rigatoni with red sauce. You will certainly dislike rigatoni with white sauce. For the purposes of this example, your enjoyment of food does not come in degrees: you either like it or not. In pictures, this is your context (the grey shaded areas represent the worlds in which you like your meal):

	red	white
spaghetti		
rigatoni		

Suppose that the contextually salient elevated priority is:

$$(18) \quad \{A \mid \text{given } A \text{ you are at least 75\% likely to enjoy your meal}\}$$

My account predicts that in the initial context the domain consists of all the *spaghetti*-worlds (the worlds that belong to the maximally ranked alternative). If, however, we acquire the information that today's sauce is red, we update to:

	red	white
spaghetti		
rigatoni		

The definition of \geq_c implies that, for a given alternative A , we only consider the portion of A that overlaps the information state supplied by c . In the example, $c + red$ forces us to compare $spaghetti \cap red$ and $rigatoni \cap red$. Only the latter satisfies the salient priority. As a result, the domain in $c + red$ consists of the worlds in which you have rigatoni with red sauce. Since clearly $spaghetti \cap red \not\subseteq (rigatoni \cap red)$, the case witnesses the expression inside the existential quantifiers in (ID).⁹

⁹Further explanation: the *spaghetti* worlds are the initial domain; the *red*-sauce worlds are the information you received. Whereas $rigatoni \cap red$ is the domain generated by (18) after updating with the information that there is red sauce.

Despite the structural similarity, then, the semantic account I just developed is not a notational variant of the classical Kratzerian semantics. My account is information dependent, and the classical account is not (Charlow, 2013, Cariani, Kaufmann and Kaufmann, forthcoming).

Let us now verify that we can solve the initial puzzle. I contrasted:

- (1) There should be beer in the fridge.
- (2) You should bet on Barcelona.

I proposed, as in Kratzer’s semantics, that the two readings of *should* arise from a single analysis by using two different priorities (more precisely, in my case: *elevated priorities*). For (1) and (2) we use, respectively:

- (19) {A | A is more likely than the alternatives}
- (20) {A | winning money is more likely than the alternatives, given A}

In a context in which (1) is true, (19) ranks as highest the alternative that beer is in the fridge. Similarly, (2) is typically true when the prejacent alternative satisfies (20) and alternatives don’t (I say ‘typically’ because there may be other, competing goals, and because a goal needn’t have the probabilizing structure of (20), e.g. (11)-(12)).

VI.

The upshot of this discussion is that, even if we are inclined to accept a probabilistic theory of *should*, the asymmetry between epistemic and deontic *should* need not be treated as a lexical difference. We can develop a premise semantics that treats it as a difference in the types of priorities that we use to rank alternatives, as long as we put our priorities in the right place. ¹⁰

¹⁰*Acknowledgments:* Thanks to Mike Caie, Dan Lassiter, Paul Portner, Aynat Rubinstein, James Shaw, Justin Snedegar, Zoltan Gendler Szabó, Malte Willer and an anonymous referee for *Thought*. This paper also benefitted from feedback I have received on related presentations at Kai von Fintel and Sabine Iatridou’s Seminar on Deontic Modality at MIT, at the Northwestern Deontic Modality Workshop and at the USC Deontic Modality Conference.

References.

- Cariani, Fabrizio (2013) "Ought and Resolution Semantics", in *Noûs*, 47(3), 534-558.
- Cariani, Fabrizio (ms.) "One Theory to Rule Them All.", manuscript, Northwestern University.
- Cariani, Fabrizio, Kaufmann, Magdalena and Kaufmann, Stefan (forthcoming) "Deliberative Modality under Epistemic Uncertainty", forthcoming in *Linguistics and Philosophy*.
- Carr, Jennifer (ms.) "Subjective Ought", manuscript, MIT.
- Charlow, Nate (2013) "What We Know and What To Do", in *Synthese*, 190, pp. 2291-2323.
- Finlay, Stephen (2010) "What Ought Probably Means and Why You Can't Detach It", *Synthese*, 177(1), 67-89.
- Horty, John F. (2001) *Agency and Deontic Logic*, Oxford, O.U.P.
- Jackson, Frank (1985) "On the Semantics and Logic of Ought", *Mind*, 94(734), 177-195.
- Jackson, Frank and Pargetter, John (1986) "Oughts, Options and Actualism", *The Philosophical Review*, 95(2), 233-255.
- Kolodny, Niko and MacFarlane, John (2010) "Ifs and Oughts", *Journal of Philosophy*, 107(3), 115-143.
- Kratzer, Angelika (1981) "The Notional Category of Modality", in B. Partee and P. Portner (eds.), *Formal Semantics: the Essential Readings*, Blackwell.
- Kratzer, Angelika (2012) *Conditionals and Modals*, Oxford, O.U.P.
- Lassiter, Daniel (2011) *Measurement and Modality*, PhD Thesis, NYU.
- Portner, Paul (2009) *Modality*, Oxford University Press.
- Sloman, Aaron (1970) "Ought and Better", *Mind*, 79(315), 385-394.

Snedegar, Justin (2012) “Contrastive Semantics for Deontic Modals”, in M. Blauw (ed.) *Contrastivism in Philosophy: New Perspectives*, Routledge, 116-133.

Williamson, Timothy (2000) *Knowledge and Its Limits*, Oxford University Press.

Windschitl, Paul and Wells, Gary (1998) “The Alternative-Outcomes Effect”, *Journal of Personality and Social Psychology* 75.6, 1411-1423.

Yalcin, Seth (2007) “Epistemic Modals”, *Mind*, 116, 983-1026.

Yalcin, Seth (2010) “Probability Operators”, *Philosophy Compass*, 5(11), 916-937.

Yalcin, Seth (2012) “Bayesian Expressivism”, *Proceedings of the Aristotelian Society*, 112(2), 123-160.