
ATTITUDES, DEONTICS AND SEMANTIC NEUTRALITY

BY

FABRIZIO CARIANI

Abstract: It has been recently suggested that a semantic theory for deontic modals should be neutral between a very large range of normative and evaluative theories. This article aims to clarify this talk of neutrality, in particular its scope and motivation. My thesis is that neutrality is best understood as an empirical thesis about a fragment of natural language that includes deontic modals – not as a new, *sui generis* methodological constraint on natural language semantics.

Introduction

It is widely agreed that there is a gap between theorizing about what one ought to do and theorizing about the meaning of deontic modals like *ought*. Similarly, there is a gap between a theory of permission and a semantic theory of *may*. In fact, this kind of gap is not peculiar to the semantics of deontic modals. We can draw parallel distinctions concerning other fragments of natural language. For example, (i) between an account of the meaning of causal language and a metaphysical account of causation (Swanson, 2010); (ii) between an account of the concept of probability and an account of the meaning of the language of uncertainty (Hamblin, 1959; Yalcin, 2010); (iii) between philosophical accounts of propositional attitudes and accounts of the semantics of attitude ascriptions (e.g. the discussion of the relationship between desire and desire talk in Fara, 2003, 2013).

Some recent works (Cariani, forthcoming; Carr, 2012; Charlow, forthcoming, all of which develop an insight of Kratzer's) take an extra step. They suggest that an adequate theory of meaning for *ought* should be neutral between a very large range of theories about what one ought to do. Similar suggestions can also be made for other deontic expressions, though

ought will remain my central example. This article aims to get clearer about how to interpret this talk of neutrality, in particular about its scope and motivation.

My thesis is that neutrality is best understood as an empirical thesis about a fragment of natural language that includes deontic modals. In particular, it is not a new, *sui generis* methodological constraint on natural language semantics and it is not best characterized as the requirement that deontic semantics avoid substantive assumptions. Instead, it is a derived constraint – one that is motivated by the same sorts of empirical considerations that guide our theorizing about modality and more generally about meaning. In adopting this position, I mean to distinguish my view of neutrality constraints from that of other proponents like Carr (2012), who thinks of it as both methodological and empirical.¹ I also mean to disagree with how some opponents like Lassiter (2014) frame neutrality constraints (as grounding non-empirical arguments).

To do so, I return to an argument for neutrality I advanced in Cariani (forthcoming) – an argument centering on the interaction between deontic modals and attitude ascriptions. I aim to integrate this argument by noting that virtually every semantic theory already incorporates a mild measure of neutrality (§2). I argue that the reasons to accept this mild degree of neutrality also support the stronger degree I advocate (§3). Furthermore, I aim to complete that argument by addressing some of the principal replies that I have encountered to it (§4).

1. Preliminaries

It will be useful to have a straightforward notation. Let A, B, C range over sentences; $\mathbf{A}, \mathbf{B}, \mathbf{C}$ range over sets of worlds. Within the same stretch of discourse and provided that A is a non-modal sentence, the value of \mathbf{A} is the set of worlds in which A is true (similarly for the other variables). Let α range over noun phrases that denote agents and φ range over verb phrases that denote actions or, more generally, properties an agent might have.

I am interested in the relationship between two kinds of theories:

- substantive theories that, given α 's deliberative situation at time t in world w , make predictions about what α ought (/is required/is allowed) to do; and
- theories of meaning for a fragment of natural language that includes deontic modals.

I call the first kind **practical theories**. I model them as (possibly partial) functions that input centered worlds (i.e. triples of the form $\langle w, t, \alpha \rangle$) and

output sets of deontic sentences (the sentences that the practical theory accepts). Because I will not be particularly interested in it, I ignore the temporal dimension of centered worlds – thus identifying them just by their w and α coordinates (the reader is invited to substitute talk of worlds with talk of world-time pairs throughout).

Some examples of practical theories are:

- a person's code of conduct (this practical theory is silent about people other than the agent whose code it is).
- the rules of basketball (this practical theory is silent about centered worlds in which the agent at the center is not playing basketball).
- a model (or a class of models) a utilitarian logic of obligation, e.g. the logics of Goble (1996) or Horty (2001, ch. 4).²

Not every function from centered worlds to sets of deontic sentences is a coherent practical theory. For example, the function that associates each centered world with every sentence of the language does not model a coherent practical theory. My argument will crucially rely on intuitions about the coherence of certain practical theories.

As for the semantics, I focus on the **compositional semantics**, a specification of the meaning of complex phrases as a function of the meanings of their constituents. This is formally captured by the interpretation function $[\cdot]$. Given a sentence A , and a sequence of parameters, $[\cdot]$ assigns to A one of two semantic values, T or F .

The sequence of parameters of evaluation includes a context c (which in the following I generally suppress), a (possibly) complex state σ ,³ and a world w . The template for this sort of account is provided by a canonical analysis of *epistemic* modals according to which they are evaluated relative to a context c , an information state i and a world of evaluation w . In the case of deontic modals, one of the central questions of the debate is what structure σ should have.

Say that a *neutrality argument* is an argument with the following structure:

- [P1] semantic theory S is incompatible with practical theory P .
- [P2] a correct semantic theory for deontic language must be compatible with P .
- [C] therefore, semantic theory S ought to be rejected.

There is significant disagreement, even among proponents of neutrality arguments, about which of them are good and why. This state of unclarity is due to (implicit) disagreements about (i) what it is for a semantic theory to be compatible with a practical theory and about (ii) which practical

theories ought to be preserved. As my argument unfolds, I will offer my general take on these questions.

2. *Warm up: mild neutrality*

We can easily imagine contexts in which (1) is acceptable.

- (1) Greta has to attend her sister's graduation.

In such contexts, the acceptability of (1) seems tied to some source – perhaps the norms that govern behavior among immediate family members, perhaps Grandpa's preferences, perhaps something else.

What is the connection between a particular source and the meaning of *have to*? The consensus view (and one of the key insights of Kratzer's work on modality) is that no particular source of requirements is encoded in the conventional meaning of deontic modals.

It might help to illustrate this idea with a toy semantics (not Kratzer's). Suppose that *have to* expresses universal quantification over a set of relevant possibilities. The set is jointly fixed by the circumstances w and by a selection function d .

$$[\alpha \text{ has to } \varphi]^{d,w} = \text{T iff for all } w' \in d(w), [\varphi(\alpha)]^{d,w'}$$

The function d models the criteria that select the permissible worlds (from the point of view of w). One set of criteria would select all and only the worlds in which Greta complies with the family norms (at w). Another would select all and only the worlds in which she complies with (a consistent fragment of) US law (in w), and so on.

This compositional semantics can be combined with radically different philosophical views about the nature of deontic discourse. According to one version of *contextualism*, d is fixed by the context of utterance; $\Gamma \alpha$ has to φ expresses in context c the proposition that is true at w iff $[[\alpha \text{ has to } \varphi]]^{d_c,w} = \text{T}$ (where d_c is c 's selection function). According to *relativism*,⁴ d is fixed by the context of assessment, and the pragmatics of deontic discourse is given in terms of the concept of truth relative to a context of utterance and a context of assessment. According to (pragmatic) *expressivism* (in the style of Yalcin, 2007, 2012), *have to*-sentences express complex contents involving d and w and the pragmatics of deontic discourse is given directly in terms of these complex contents.

The toy semantics for *have to* presupposes a weak neutrality constraint. As long as the prejacent is consistent, the semantics does not predict *a priori* that any *have to*-sentences are acceptable (or unacceptable, as the case might be). For instance, the semantics does not privilege 'You have to

be polite' over 'You have to be rude'. Our general inclination to accept the first statement and reject the second must be partly explained by non-semantic facts.⁵ In general, we might require:

Mild Neutrality: for every coherent practical theory P and centered world $\langle w, \alpha \rangle$, if $\Gamma \alpha$ has to $\varphi \top \in P(\langle w, \alpha \rangle)$, then there is a selection function d such that $[\alpha$ has to $\varphi]^{d,w} = \text{T}$.

Informally, this says that any single *have to*-sentence that is accepted by a practical theory can be captured by some selection function or other. Having stated Mild Neutrality, we should ask: what justifies it?

In my view, one of the central arguments in its favor is that restricting the admissible practical theories leaves us unable to predict how *have to*-sentences embed in attitude contexts. Imagine that there was such a restriction – for example, that there was no d and w such that you behave rudely at every world in $d(w)$. Suppose also that David belongs to a cult that values rudeness above all else. We should be able to say:

(2) David thinks that you have to be rude (but in fact you don't).

If no selection functions encoded rudeness-promoting practical theories, it would be difficult, if not downright impossible, to make sense of (2). One could, perhaps, try to claim 'you have to be rude' is a contradiction and so (2) is handled by whatever devices handle ascriptions of inconsistent beliefs. However, claiming that 'you have to be rude' is a contradiction is not particularly plausible.

This insight can be made a bit more precise if we combine the toy semantics with an account of attitude verbs. According to a prominent tradition (pioneered by Hintikka, 1962), attitude verbs are interpreted as quantifiers over possibilities. The plainest such account is:

$[\alpha$ thinks that $A]^{d,w} = \text{T}$ iff for all worlds w' compatible with α 's state in w , $[A]^{d,w'} = \text{T}$

In the case of beliefs with modal content, it is plausible (following Stephenson, 2007, and Yalcin, 2007, 2012) to adopt a version of the operator analysis that shifts the relevant state:

$[\alpha$ thinks that $A]^{d,w} = \text{T}$ iff for all worlds w' and selection functions d' compatible with α 's state in w , $[A]^{d',w'} = \text{T}$

The idea is that we do not just shift the circumstances to match α 's belief worlds but also the criteria for permissibility encoded by the selection function.

The combined predictions of this analysis of *thinks* and the toy semantics for *have to* are:

$[\alpha \text{ thinks that } \beta \text{ has to } \varphi]^{d,w} = \text{T}$ iff for all worlds w' and selection functions d' compatible with α 's state in w , for all $w'' \in d'(w')$, $[\varphi(\beta)]^{d',w''} = \text{T}$.

The argument for Mild Neutrality is that a deviant, rudeness-promoting selection function must be available for the theory to make the right prediction about (2).

A possible concern about this argument is that there are notorious problems with operator accounts of *thinks*. The concern is important and worth addressing (I will do so in §4). The present point, however, is just that an informal argument for Mild Neutrality can be made precise under one of the most prominent hypotheses about the behavior of attitude verbs.

In fact, this style of argument has a precedent in the epistemic modals literature. Consider the proposal that $\ulcorner \text{It might be the case that } A \urcorner$ means, roughly, that A is compatible with the beliefs of a contextually salient group. This proposal is refuted by the observation that, when embedded under *thinks*, epistemic *might* gets interpreted relative to the state of the subject of the ascription (the point is made in Egan *et al.*, 2005; Stephenson, 2007; Ninan, 2010; the example is from Stephenson).

(3) Sam thinks it might be raining.

The take-home point is that individual epistemic states have to be accessible to the semantics (alongside group ones), in order to make sense of (3). The case of deontic *have to* seems parallel: in deontic attitude ascriptions we want to shift the state to match the state of the subject of the ascription. But this requires the semantic apparatus to offer a rich enough menu of possible states.

These considerations can be extended in two ways. Given that none of the arguments turn on specific features of *have to*, they immediately extend to other deontic modals, like *ought*, *should* and *may*.

The second extension is that the arguments can be applied (with appropriate modifications) to more complex semantic analyses. Consider an ordering semantics that blends some ideas of Lewis (1973) with some of the ideas of Kratzer (1977, 1981, 2012). On this analysis, *ought*-sentences are evaluated relative to a modal background i (a set of worlds) and a function $<$ that assigns to each world w a pre-order of worlds $<_w$. Instead of encoding criteria for permissibility, $<$ captures the comparative goodness of worlds (according to some standard or other). The lower a world is in the ordering the better (an image that is frequently used is that the world is closer to the contextual ideal of goodness). On the simplifying assumption that (for every w) $<_w$ only draws finitely many distinctions,⁶ the semantics can be stated as:

$[\text{ought: } B]^{i,<,w} = \text{T}$ iff for all worlds w' that are $<_w$ -minimal in i , $[B]^{i,<,w'} = \text{T}$

A world w' is $<_w$ -minimal in i , if w' belongs to i , and no other worlds in i are better than w' (i.e. ranked below w'). Like the earlier semantics for *have to*, this analysis satisfies the appropriate version of Mild Neutrality:

Mild Neutrality: for every coherent practical theory P and centered world $\langle w, \alpha \rangle$, if $\Gamma \alpha$ ought to $\varphi \top \in P(\langle w, \alpha \rangle)$, then there are parameters $i, <$, such that $\llbracket \alpha$ ought to $\varphi \rrbracket^{i, <, w} = \text{T}$.

Again, the idea is that it is not up to the semantics to constrain what one ought to do, any more that it is up to the semantics to constrain which cars can be red. If we build a particular conception of comparative goodness of worlds into the semantics, we end up unable to model embeddings under attitude verbs. Instead, if we leave the possible values of i and $<$ totally unconstrained, we can combine the semantics with a state-shifting operator analysis and get the desired flexibility.⁷

3. *Generalized neutrality*

Very few interesting theories are ruled out by Mild Neutrality. In this section, I formulate and defend a generalization of Mild Neutrality that is incompatible with genuinely interesting approaches.

Mild Neutrality is based on single deontic sentences. This raises the question: are there empirically motivated neutrality constraints that apply to *sets* of sentences? Obviously, we cannot demand that for *any* set of deontic sentences there be a parameter assignment relative to which they come out all true. After all, some sets of deontic sentences are straightforwardly inconsistent. However, we can require that, for every set of sentences that are jointly accepted by a coherent practical theory, there be an eligible assignment to the parameters in the semantics relative to which all of those sentences come out true.

Formally, the appropriate generalization of Mild Neutrality is:

Generalized Neutrality: for every coherent practical theory P there is a state σ such that for every centered world $\langle w, \alpha \rangle$, for every sentence A if $A \in P(\langle w, \alpha \rangle)$, $\llbracket A \rrbracket^{\sigma, w} = \text{T}$.

Though they look similar, Generalized Neutrality is much stronger than Mild Neutrality. Because the quantification over states σ is made independently of the choice of centered worlds and the choice of sentence (in logical terms, the quantifier order of Mild Neutrality is $\forall \forall \forall \exists$; the quantifier order in Generalized Neutrality is $\forall \exists \forall \forall$).

I propose that Generalized Neutrality is justified by the same empirical considerations involving attitudes that justify Mild Neutrality. It is this generalized notion of neutrality that can be put to serious use in interesting neutrality arguments. This section has two parts: §3.1 sketches the account

of the meaning of deontic *ought* that my argument targets; §3.2 develops a specific neutrality objection against this account.

3.1. THE TARGET

A prominent target for neutrality arguments are views that explain the meaning of deontic *ought* (*should*, and other deontics) in terms of the mathematical notion of expected value. One of the key insights of these views is that a system based on orderings, such as the Lewis/Kratzer account, misses out on quantitative comparisons in the evaluation of worlds.⁸

The proposed remedy is to include among the parameters of evaluation a numerical assignment of the value to the individual worlds. The formal tool that captures this assignment is a *value function*. A rough take on the compositional semantics is that Γ *ought*: $A \sqsupset$ is true (relative to appropriate parameters) just in case the expected value of **A** is greater than the expected value of the alternatives. On one way of calculating it, the expected value of **A** is the weighted average of the value of the worlds in **A**, where the weight of each world w is $Pr(\{w\} \mid \mathbf{A})$.⁹ At a minimum, the state σ that is needed for an expected value semantics must include a modal background i (a set of worlds); a probability function Pr ; a set of alternatives Alt and a value function v . When these parameters are given, I write ' $Ev(\mathbf{A})$ ' to denote the expected value of **A**.

The rough characterization of the previous paragraph deploys a deliberately ambiguous phrase ('greater than the expected value of the alternatives'). Both disambiguations of this phrase give rise to possible implementations of the view. On one disambiguation, Γ *ought*: $A \sqsupset$ requires that $Ev(\mathbf{A}) > Ev(\mathbf{B})$ for each alternative **B**. On the other, the weaker condition $Ev(\mathbf{A}) > Ev(\mathbf{B}_1 \cup \dots \cup \mathbf{B}_k)$ is enough.¹⁰

Developing an expected value account requires decisions at several other choice points. Two of the most important are:

- (I) *Does the value function for α at t in w coincide with α 's subjective utility function (at t in w) in the decision theorist's sense?*

The consensus is that it does not. The value function v should be allowed to represent many other ways of assigning value to worlds. For example, there should be value functions that track how α objectively ought to evaluate individual worlds, independently of the subjective desirability α assigns to those worlds.

One reason in favor of the consensus opinion is that we want to satisfy Mild Neutrality. If the value function were identified with α 's subjective utility function, we would not be able to express the oughts that correspond to these other forms of evaluation.

Another reason is an empirical argument proposed by Yalcin (2012, p. 146). Yalcin argues that identifying the value function with the agent's subjective utility function predicts unexpected logical connections between belief ascriptions in deontic contents ('Judy thinks Vann ought to apologize') and desire ascriptions ('Judy wants Vann to apologize') – assuming that desire ascriptions also go in for an expected value analysis. The second important choice point is:

(II) *what is the proper logic for the expected value ought?*

The expected value accounts canvassed so far project a radically non-classical deontic logic. However, as Wedgwood (forthcoming) notes, it is easy to give a semantics based on expected values that outputs standard deontic logic (if one wants to). All we have to do is use expected values to rank alternatives and then say that *ought* universally quantifies over the worlds that belong to the maximal alternatives.¹¹ My argument in §3.2 targets expected value views regardless of how they fall on this choice point.

3.2. THE PROBLEM

My objection can be developed with almost any textbook decision problem. Because it has been so familiar in recent literature, I will use a variant on the much-discussed miners paradox of Regan (1980).¹²

Biased Coins: Joan offers Kate a bet on the outcome of a 100% biased coin. The coin's bias is unknown. The bet costs \$1 and pays \$1.90 if Kate's prediction is correct. Kate values only acquiring money and values each dollar equally. Kate can also refrain from betting.

The standard miners paradox turns on judgments about claims like:

- (4) Kate ought to refrain from betting.
- (5) Either Kate ought to bet on heads or Kate ought to bet on tails.
- (6) If the coin is biased towards heads, Kate ought to bet on heads.

In the present context, I want to bracket the question whether or not we judge these claims to be true. For what it's worth, my view is that (a) the provided context is not described in enough detail to yield determinate judgments about (4) or (5), but (b) on the natural way of filling it out, I am inclined to accept (4) and (6) and reject (5).

My interest here is in data about belief ascriptions involving the contents of (4) and (5). To get at those belief ascriptions I will describe the states of two fictional individuals, Average and Cautious.

- Average accepts (4) and rejects (5). His thought process is to ascribe a probability to each bias (let's stipulate: 50/50), set up a proper

decision problem and rank Kate's alternatives (i.e. bet on heads, bet on tails, and refrain) according to their expected value. Average takes Kate's value function to be proportional to Kate's payoffs.

- Cautious rejects both (4) and (5). He thinks that, when the probabilistic information reflects total ignorance (rather than evidence about chances), Kate's alternatives ought to be ordered by a dominance criterion. In the relevant sense, **A** weakly dominates **B** relative to background i iff every world in $(\mathbf{A} \cap i)$ is at least as good as any world in $(\mathbf{B} \cap i)$, and some world in $(\mathbf{A} \cap i)$ is better than any world in $(\mathbf{B} \cap i)$.

Obviously, Average and Cautious do not represent the only possible takes on the example. A believer who followed the actualism of Jackson and Pargetter (1986) would end up in a rather different place. Under actualism, the ranking of an alternative **A** is determined by the value of the most similar **A**-world to the actual world. In **Biased Coins**, actualism invites us to compare the values of the closest world in which Kate refrains with the value of the closest world in which Kate takes one of the bets. The actualist believer would accept (5) and reject (4) (which makes sense, since actualism is intended to be an analysis of the objective *ought*).

Having thus described the states of Average and Cautious, we should be able to make the following belief ascriptions:

- (7) Average thinks Kate ought to refrain from betting.
- (8) Cautious thinks it's false that Kate ought to refrain and thinks it's false that Kate ought to bet.

My objection is that the expected value semantics runs into trouble with (8). More generally, that it cannot properly model the attitudes of agents who do not factor expectations in how they rank alternatives. The intuitive problem with (8) is that Cautious does not have beliefs whose content are well modeled as comparisons of expected values.

In Cariani (forthcoming), I made this objection more precise by assuming the state-shifting operator analysis of *thinks*. In the remainder of this section, I repeat that part of the argument. In the next, I consider how it extends beyond such theories. The state-shifting analysis of *thinks* that fits the expected value approach to deontic modality is:

$[\alpha \text{ thinks that } B]^{i,Pr,v,Alt,w} = \text{T}$ iff for all i', Pr', Alt' and v' compatible with α 's state, and for all $w' \in i'$, $[B]^{i',Pr',v',Alt',w'}$.

Because Average and Cautious must end up having conflicting attitudes, there must be a difference in at least one of the i', Pr', Alt' and v' that are compatible with their states.

This difference cannot come from either i' , Pr' or Alt' . This is because we can just stipulate that Average and Cautious have the same information, both qualitative and probabilistic. We can also stipulate that they agree on Kate's available options. It follows that any difference between them must be traced to a difference in v' , their value function.

To derive (7), it is enough to suppose that the value functions compatible with Average's state assign value in proportion to Kate's monetary payoff. For instance, if in w Kate does not bet, $v(w) = 0$; if she bets and wins, $v(w) = +0.9$; if she bets and loses $v(w) = -1$. Then, the expected value of refraining is 0, while the expected value of taking either bet is negative ($-1 \cdot .5 + .9 \cdot .5 = -0.05$, assuming that probabilities for each bias are .5). This would match Average's beliefs as described in (I).

The situation is trickier when it comes to figuring out the value function that is compatible with Cautious's state. Here the expected value theorist does not have the luxury of choosing a value function that makes independent sense. The value function must be reverse-engineered from the verdicts we hope to derive. Because the expected value account does not allow for incomparabilities, the only way to derive Cautious's judgments is if Cautious's expected values are all identical, i.e.:¹³

$$(*) \quad Ev(\mathbf{refrain}) = Ev(\mathbf{bet-heads}) = Ev(\mathbf{bet-tails})$$

This assignment, however, is problematic. It will turn out that this representation of Cautious's state does not correctly model the rest of his beliefs, in particular his beliefs in conditional *ought* sentences.

To see this, let's add to the story: suppose that there is a possible piece of evidence that strongly predicts that the coin is biased towards heads. Perhaps, if Joan smirks, the coin is 90% likely to be biased towards tails. Now consider:

- (9) Cautious thinks it's false that if Joan smirks, Kate should bet on heads.

Given Cautious's state, (9) should be true. Cautious ranks alternatives by a dominance criterion: even supposing Joan smirks, betting on heads does not dominate the alternatives.

However, the expected value theory cannot make this prediction. As the probability of heads increases, so should the expected value of **bet-heads**. As recorded in (*), Kate's three options were tied in expected value according to Cautious's state. It is a simple consequence that if the probability of heads increases, the expected value of betting on heads should increase accordingly. In slightly more general terms, the problem is that twisting Cautious's v -function so as to make it fit (8) makes it impossible to make it fit (9).

This completes my presentation of the objection against expected value accounts in its basic version. Summing up, we started out by combining an expected value semantics with the operator analysis for *thinks*. Given this, I asked how we should model the beliefs of a non-Bayesian believer (in our case: Cautious). If Average and Cautious share a value function, (8) is not predicted. To predict it, we must ascribe to Cautious a different value function, one that ensures (*). But the only value function that works (given the story in **Biased Coins**) is one that fails to predict (9).

At this point, we can run a neutrality argument. The expected value semantics is incompatible with Cautious's practical theory. But Cautious's practical theory seems coherent, even if one disagrees with it (as I do). The theory of meaning should not *a priori* rule out Cautious's practical theory any more that it should determine that you have to be polite rather than rude. In particular, if it did rule out Cautious's practical theory, we would fail to make sense of ascriptions of deontic beliefs to him. If we adopt Generalized Neutrality, we should reject the expected value account.

4. *Objections and replies*

OBJECTION 1: THE ARGUMENT DEPENDS ON THE OPERATOR ANALYSIS OF *THINKS*. BUT THIS ANALYSIS IS NOTORIOUSLY PROBLEMATIC

Fine, the analysis is problematic. However, my argument does not really depend on its problematic features. To show this, I will explore how the argument plays out on the principal rival of the operator analysis: the view that *thinks* expresses a relation between subjects and contents. Call this the *relational analysis*. Unlike the operator analysis, it is hard to trace the relational analysis back to a first pioneer (though Russell is a good candidate).

Here I follow one recent implementation of the relational analysis (MacFarlane, 2014, p. 154).

$$\begin{aligned} [\textit{thinks}]^{c,w,t} &= \{ \langle x, y \rangle \mid x \text{ has a belief with content } y \text{ at world } w \text{ and time } t \} \\ [\textit{that } A]^{c,w,t} &= |A|^c \end{aligned}$$

In MacFarlane's system $|A|^c$ denotes the content of A in context c (the system is neutral about the underlying theory of contents). MacFarlane gives two reasons to prefer this account to the operator analysis: (i) it makes sense of sentences like 'There is something that both Joe and Mary are thinking'¹⁴ and (ii) it does not imply the closure of *thinks* under necessary entailment. The first problem is clearly of no concern here: all the attitude reports I have considered have sentential complements.

Neither does the objection of §3.2 appeal to any principles of closure of belief under necessary entailment. Recall that the objection relies on these premises:

(judgment 1) Both (8) and (9) ought to be predicted true.

(logic) either the contextual parameters satisfy claim (*) or they do not.

(prediction 1a) If they do, (9) is predicted false.

(prediction 1b) If they do not, (8) is predicted false.

The reasoning that supports these premises does not exploit closure under single-premise necessary entailment (when $\mathbf{A} \subseteq \mathbf{B}$, $\Gamma\alpha$ thinks that $A \sqsupset$ entails $\Gamma\alpha$ thinks that $B \sqsupset$) nor multi-premise closure principles (e.g., in the case of conjunction introduction, $\Gamma\alpha$ thinks that $A \sqsupset$ and $\Gamma\alpha$ thinks that $B \sqsupset$ entails $\Gamma\alpha$ thinks that $A \& B \sqsupset$).

Despite avoiding the use of closure principles, my reply is not complete. One might still object that underlying predictions 1a and 1b is a different bad feature of the operator analysis. To be specific, the problem arises because the operator analysis does not allow agents to have inconsistent beliefs. This is relevant because, as we will see shortly, the contents of the beliefs ascribed in (8) and (9) turn out to be inconsistent by the lights of the expected value analysis (in a somewhat non-standard sense of ‘inconsistent’). As a reminder, the contents of the beliefs ascribed in (8) and (9) are:

- (i) that it is false that Kate ought to bet
- (ii) that it is false that Kate ought to refrain and
- (iii) that it is false that if Joan smirks, Kate ought to bet on heads.

According to the expected value analysis, there is no single parameter assignment that vindicates (i)-(iii) and is faithful to what is stipulated in the case. As we will see, I disagree with the assessment of the significance of this result: far from rescuing the expected value analysis, the result points to a way of setting up my objection without relying on the operator analysis at all.

But let’s slow down: say that a parameter assignment is *faithful* to a background story (like the text of **Biased Coins**) just in case every constraint on the parameters that is specified in the text is reflected in an appropriate restriction on the value of the parameters. If the text claims that some proposition **A** has probability .73, a faithful assignment’s probability coordinate *Pr* ought to assign $Pr(\mathbf{A}) = .73$. If the text claims that an agent has two alternatives, say raising her left arm or raising her right arm, then a faithful assignment must restrict *Alt* to **{raise left, raise right}**.

The story of **Biased Coins** states that Kate has exactly three options: refraining from betting, betting on heads and betting on tails. Furthermore, when I introduced (9), I stipulated: $Pr(\text{the coin is biased towards$

heads| Joan smirks)=.9. We could also add to the story (without affecting any judgments) that this is significantly higher than $Pr(\text{the coin is biased towards heads})$ – indeed, in making calculations, I let this be .5. These stipulations are sufficient to support:

Key Fact 1: There is no parameter assignment $\langle i, Pr, v, Alt, w \rangle$ that is faithful to **Biased Coins** and such that (i), (ii) and (iii) are all true.

The argument of §3.2 doubles up as a proof of Key Fact 1: the only faithful parameter assignments that verify (i) and (ii) give equal expected value to the three alternatives. But that won't work for (iii). Key Fact 1 shows that, regardless of our favorite analysis of propositional attitude reports, (8) and (9) ascribe inconsistent beliefs to Cautious.

It is at this point that I part ways with defenders of the expected value analysis. In particular, we disagree on:

(judgment 2) the contents of Cautious's beliefs ought to be jointly consistent given the background story.

If judgment 2 is correct, so that (i)-(iii) are consistent, the expected value analysis is in trouble (unless supplemented by another explanation for the apparent consistency).¹⁵ If judgment 2 fails, so that (i)-(iii) are inconsistent, any theory that predicts their consistency is in trouble (unless supplemented by another explanation of the apparent inconsistency).

I grant that it is reasonable to be less confident in judgment 2 than in judgment 1. However, on balance I believe that there is a stronger inclination to view the verdicts entailed by dominance orderings (or actualist orderings, or *maximin* orderings, or minimal variations on the particular expected value ordering we characterized) as consistent than to treat them as inconsistent.¹⁶ I am unmoved by claims to the effect that we should privilege the 'true theory' of rational choice. For one thing, the expected value analyses that are presented in semantics are almost certainly not the true theory. For another, no one appeals to the 'true theory' at the level of Mild Neutrality. There is no reason to think the present case different.

Accepting judgment 2 does not force us to treat all kinds of unorthodox practical theories as consistent. As I will highlight in my replies to Objections 2 and especially Objection 5, there might well be interesting and substantive practical theories that nonetheless are incoherent and give rise to inconsistent judgments. If so, there is no demand on the semantics to make those judgments consistent. This point allows me to emphasize that all that is needed for my argument is that *some* practical theories *sometimes* make consistent predictions that are incompatible with the predictions of the expected value account. I supposed that dominance-based practical theories and **Biased Coins** are a good choice, but the particular choice of example is inessential if the general point is granted.

One last point: whether or not the reader shares judgment 2, I hope to have delivered on one of the promises I made in the Introduction – to vindicate the claim that the dispute on semantic neutrality is fundamentally empirical. Just like many other disputes in semantics, it bottoms out in judgments about consistency and entailment.

OBJECTION 2: THE ARGUMENT PROVES TOO MUCH

Let me give this objector a bit more room to articulate her concern:

An analogue of the argument of §3.2 could be applied against probabilistic semantic theories for probability operators. For example, on a rudimentary probabilistic semantics, [*probably*: A]^{Pr,w} iff $Pr(A) > .5$. This builds substantive assumptions about probability into the meaning of *probably*. Imagine an agent who believes that it is not likely that it will either rain or not, or perhaps an agent that ascribes greater probability to a conjunction than she does to one of its conjuncts. If the argument of §3.2 was successful, it would seem to refute any semantics whose logical truths include ‘It is likely that it will either rain or not’ or $\Box A$ is at least as likely as $(A \ \& \ B)\Box$. But that is utterly implausible.

My reply to this objection is twofold. First off, it is not totally implausible to attack the probabilistic semantics for validating too many inferences. Even theorists who favor probabilistic accounts sometimes back off to weaker systems (for example, by retreating to sets of probability measures). So, one should not quickly concede that anything that is validated by the probabilistic semantics is a *bona-fide* validity.

Having said that, it seems sensible to think that ‘rain and wind is not more likely than rain’ should be valid on any semantics, probabilistic or not. This plausible idea seems in conflict with the fact that the following belief ascription seems perfectly acceptable (provided that Walter is suitably irrational).

(10) Walter thinks that rain and wind is more likely than rain.

It is implausible to conclude from the acceptability of (10) that we should adopt a semantics according to which ‘rain and wind is more likely than rain’ is consistent.

I agree that this consequence is implausible. However, there are significant differences between (10) and Cautious’s trio of beliefs. The difference is that (10) ascribes to Walter a belief in a content that we clearly judge inconsistent. As I pointed out, I don’t think there is nearly as strong a judgment that Cautious’s beliefs are inconsistent (though, again, I do disagree with them).

Given this, (10) ought to be treated by whatever medicine can cure the problem of logical omniscience. But it does not follow that we should do the same with Cautious’s beliefs.

One last point: I have scrupulously avoided speaking of ‘building substantive assumptions’ in constructing my argument. The force of neutrality arguments does not arise from refraining from ‘building in substantive assumptions’ (this might be impossible). Rather, it comes from a need to not rule out coherent, if not always plausible, ways of evaluating alternatives.

OBJECTION 3: THE ARGUMENT PROVES TOO MUCH (AGAIN)

The objector continues:

Still there are belief ascriptions that do not trigger judgments of inconsistency and in which we do not want to go for parallel neutrality constraints. Consider disputes over the extension of scientific terms, such as the famous dispute about whether Pluto is a planet. Consider two scientists (call them *In* and *Out*). According to *In*’s usage, Pluto counts as a planet. According to *Out*’s usage, to be a planet a body needs to clear its neighborhood from debris. Suppose that it is a metasemantic fact that *Out*’s usage is correct. Nonetheless, *Out* can plausibly say:

(11) *In* thinks that Pluto is a planet.

In asserting (11), *Out* does not ascribe to *In* the belief that Pluto cleared its orbit from debris. Instead, the belief that *Out* ascribes to *In* seems to be that Pluto counts as a planet – in light of *In*’s (deviant) criteria. Somehow, attitude ascriptions can latch on the meanings of deviant speakers.

My reply is that this case too is relevantly disanalogous from the neutrality argument of §3. Specifically, I never claimed that Average and Cautious have different metalinguistic beliefs about the meaning of *ought*. I just required that they rank alternatives according to different and incompatible criteria (in my language: that they embrace incompatible practical theories).

It might help to remember here that Cautious’s beliefs are meant to pattern with the beliefs of David (the member of the rudeness-promoting cult from §2). There is no temptation, I think, to say that David uses *have to* differently from the rest of us.

I grant that there are sophisticated moves one might want to make to handle (11) (allow for metalinguistic belief ascriptions, go for analogues of Stalnaker’s diagonalization’s strategy, etc.). But it is implausible to claim that these maneuvers ought to carry over to case of deontic belief ascriptions.

OBJECTION 4: A SLIGHT MODIFICATION OF THE EXPECTED VALUE ACCOUNT CAN GET THE CONSISTENCY VERDICT

The semantics of Lassiter (2011) has more resources than I have suggested. In particular, he notes (§3.6.3) that even the expected value semantics can make room for incomparability.

As van Rooij [. . .] points out with respect to adjectives like *clever*, it is possible to capture incomparability without weakening scales too much by treating a scale as a set of subscales. Each of these is built around a connected order, and the global truth-conditions of $x \geq_p y$ rely on universal quantification over all the value of $x \geq_p y$ [. . .]. So, for example, the semantics would treat *John is cleverer than Mary* as true if John is cleverer than Mary with respect to every subscale [. . .], false if John is not cleverer than Mary with respect to any subscale [. . .] and undefined otherwise (Lassiter, 2011, p. 81).

Translating to the present case, instead of ascribing to Cautious a single value function v , we might ascribe to him a set of value functions $\{v_1, \dots, v_k\}$. We could then say that an *ought*-sentence is true just in case the ordinary expected value semantics predicts its truth at each v_x in the set.

Perhaps, then, Cautious's state is a pair of value functions, such that relative to one function Kate ought to refrain and relative to the other Kate ought to bet. If that's how we should represent Cautious's state, we can indeed predict that Cautious would not believe either (4) nor (5).

There are a few reasons to be skeptical of this strategy (to be clear: Lassiter did not suggest it either in print or in conversation as an answer to the problem of §3.2). First, it is not clear what are the two incomparable dimensions that represent Cautious's state. Second, the problem of §3.2 can be replicated with any deviant decision rule, while the present answer only affects the case of someone who ranks alternatives by dominance. In Cariani (forthcoming), the role of Cautious was played by a believer in *Maximin*, a non-probabilistic decision rule that does not allow for incomparabilities. Third, though this approach might work for (8)-(9), there are suppositions under which Cautious no longer thinks that Kate's options are incomparable, e.g.:

- (12) If the coin is biased towards heads, Kate ought to bet on heads.

If the 'pair representation' of Cautious's state derives (9), it is unclear how it can also derive the claim that Cautious believes the content of (12).

OBJECTION 5: EVERY SEMANTIC THEORY EITHER ALLOWS OR DISALLOWS MORAL DILEMMAS. EITHER WAY, THE THEORY MUST MAKE A SUBSTANTIVE ASSUMPTION. SO, EVERY SEMANTIC THEORY MUST VIOLATE SOME NEUTRALITY CONSTRAINTS

The case of moral dilemmas is complex enough to deserve its own separate treatment. I limit myself to some summary observations that (I hope) take the sting out of this objection.

As should be clear by now, I am not claiming that the semantics should be free of every substantive assumption. I have taken the rather different approach of presenting neutrality as the requirement that the semantics be compatible with every coherent practical theory.

This difference in stance helps with the case of moral dilemmas. I think it is fair to say that the dispute about whether there are or are not moral dilemmas concerns the very coherence of a practical theory that allows for them.

If moral dilemmas are incoherent, then the semantic theory should ban them. Belief in moral dilemmas would then be treated in the same way as belief in 'rain and wind is more likely than rain' (as discussed under Objection 2).

If moral dilemmas are coherent, the semantic theory should access states that support them. There is no violation of Neutrality in that: within a semantics that contains states that support moral dilemmas, there will also be states that do not support them. So, it is possible to model the beliefs of those who reject moral dilemmas.¹⁷

Department of Philosophy
Northwestern University

NOTES

¹ On the one hand, Carr develops a convincing empirical objection against the proposal of Cariani *et al.* (2013). On the other, she seems to suggest that one can determine abstractly whether certain semantic theories 'make substantive assumptions', which she treats as inherently objectionable. I should add here that Carr's (2012) is a preliminary version of a longer, currently unpublished paper titled 'Subjective Ought' that has influenced my thinking on these matters (though I disagree with a number of points in it).

² The reference to a class of models is essential because the logic itself, in abstraction from any models, does not have significant implications about what one ought to do.

³ The nature of σ varies from theory to theory. On some theories, it is a pair consisting of an information state and a selection function. On others, it is a pair consisting of a probability function and a utility function. Many other options are possible. Instead of trying to come with a clumsy description that holds together all the possible ways of relativizing the analysis of deontic modals to parameters, I invite the reader to think of it as 'whatever parameters in addition to c and w are necessary to specify the interpretation of deontic modals'.

⁴ This form of relativism is in the style of MacFarlane (2014). However MacFarlane does not think that 'normative' parameters in the semantics of deontic expressions (like d in the toy semantics for have to) are determined by the context of assessment.

⁵ Perhaps the explanation is that conversational participants generally have the sort of attitudes that determine domains that are incompatible with rudeness and entail politeness, together with metasemantic principles that connect these attitudes with the parameters in the semantics.

⁶ Both Lewis and Kratzer officially reject this assumption. In fact, they reject the much weaker Limit Assumption. The limit assumption is the claim that the set of $<_w$ -minimal worlds is non-empty. As noted in (Cariani *et al.*, 2013, fn 10), the kind of lexical entry I

provide here requires a strengthening of the Limit Assumption to the claim that every linearly ordered sub-chain within $<_w$ terminates in a minimal element. To say that $<_w$ draws only finitely many distinctions asymmetrically entails this more complex condition.

⁷ I note in passing that Lewis seems to have endorsed a version of Mild Neutrality (though for somewhat different reasons). The issue is first discussed in Counterfactuals:

As is the custom in deontic logic, I shall say nothing definite about the source and significance of this ordering. Perhaps the worlds are ordered according to their total net content of pleasure, measured by some hedonic calculus; or their content of beauty, truth and love; or their content of some simple, non-natural quality. Perhaps they are ordered according to the extent that their inhabitants obey the law of God, of Nature or of man. Perhaps according to how well they measure up to some sort of standards of objective morality, if such there be; perhaps according to someone's personal taste in possible worlds; perhaps according to calm, sympathetic, impartial contemplation of alternative possibilities. It does not matter. We can build in the same way on any of these foundations, or on others (Lewis, 1973, p. 96).

Here Lewis appears to be simply reporting on (and following) a custom of distinguishing two stages in constructing a semantics. But, by the time the point is taken up in his 'Reply to McMichael' (Lewis, 1978) his stance is more clearly an endorsement of the division of labor captured by Mild Neutrality.

The semantic analysis tells us what is true (at a world) under an ordering. It modestly declines to choose the proper ordering. That is work for a moralist, not a semanticist.

Still, as I read the quote there remains a difference between Lewis's position and mine. Lewis does not exclude the possibility that, say, *ought* might end up meaning *ought*-relative-to-source-s. Lewis's view seems to be that the deontic semantics is best practiced as a modular enterprise, in the sense that it should fit any foundation we might decide to supply for it.

⁸ Among the central predecessors of this view are Prior (1956); Jackson (1991); and Goble (1996). One important caveat, however, is that these works are not best interpreted as developing a linguistic picture of the meaning of deontic modals. As I noted in §1, Goble provides a formal analysis of a utilitarian obligation operator. Similarly, Jackson's project is to characterize a probabilistic consequentialism that can address some classical objections against non-probabilistic consequentialist views. For works in which the expected value analysis is discussed as a linguistic hypothesis, see Cariani, 2009; Lassiter, 2011, 2014; Yalcin, 2012; Wedgwood, forthcoming.

⁹ This is the way of Lassiter (2011, §6.3.3). Wedgwood (forthcoming, §3) gives a subtler, but more complex definition that is closer to what is standard in decision theory. The differences do not matter here, so I go with the simpler option.

¹⁰ The former is proposed in Cariani, 2009. The latter is tentatively suggested in Lassiter, 2011. The two positions collapse if we suppose that the set of alternatives to **A** is $\{\bar{\mathbf{A}}\}$. If we think that ought-sentences can genuinely express a comparison among multiple alternatives, the former proposal seems more promising, because $Ev(\mathbf{A}) > Ev(\mathbf{B}_1 \cup \dots \cup \mathbf{B}_k)$ is compatible with the possibility that $Ev(\mathbf{B}_i)$ is much greater than $Ev(\mathbf{A})$, in which case, I think we would want to reject $\Gamma \text{ought: } \mathbf{A} \top$.

¹¹ This technique for constructing a quantificational domain out of an ordering of alternatives is familiar independently of expected value theories. Horty (2001) uses it in the context of a dominance-based deontic logic to construct a domain out of a ranking of alternatives (see his §4.2.1 for Horty's discussion of how to extract a set of optimal actions out of a dominance ordering and his §4.3.1 for his definition of a deontic necessity operator).

¹² In recent literature, the miners paradox has been used as part of an attack on *modus ponens* for deontic conditionals (Kolodny and MacFarlane, 2010; see Willer, 2012 for a reply), as well as an challenge against theories like the Lewis-Kratzer semantics I sketched in the previous section (see Cariani *et al.*, 2013; Charlow, 2013 for the challenge, Dowell, 2012; Bronfman and Dowell, forthcoming; von Fintel, unpublished, for a defense of the Lewis-Kratzer approach). Here I want to set aside both debates, and use the example for rather different purposes.

¹³ Strictly speaking this is not true in Lassiter's theory for two reasons. (a) For Lassiter, Γ ought: $A \uparrow$ is true (relative to appropriate parameters) iff the expected value of A significantly exceeds the expected value of the alternatives. So Cautious's state might have differences in expected value provided they are not very significant; (b) Lassiter does have a way of introducing a modicum of incomparability in his system. Neither of these points affects my argument: there is an obvious way of recasting my claims in this section to work around (a). As for (b), it is addressed as Objection 4 in the next section.

¹⁴ MacFarlane actually discusses *believes*, and notes more plausible instances of (i), like 'Joe believes Goldbach's Conjecture'. Such examples are not so smooth with *thinks* in place of *believes*.

¹⁵ I say 'in trouble' rather than 'refuted' because none of these arguments can reasonably amount to a refutation of the expected value semantics. It is possible, for all I have said, that the considerations that favor the expected value semantics far outweigh the costs I am outlining here. Having said that, the project of Cariani (forthcoming) is precisely to show that it is possible to design a semantics that satisfies General Neutrality and yet shares the principal advantages of the expected value semantics.

¹⁶ One could even go a step further and agree with Buchak's (2014) view that some of these orderings should not even be deemed irrational from the decision theoretic point of view, but I will remain neutral on this additional step.

¹⁷ Thanks to Steve Finlay, John MacFarlane, Mark Schroeder and Tim Sundell for feedback on a draft of this article. The article is part of a larger project of which Cariani (forthcoming) is the centerpiece. Consequently, it has benefitted from feedback I have received as part of that project. In particular, from conversations with Nate Charlow, Steve Kuhn, Daniel Lassiter, Shyam Nair, Paolo Santorio and Malte Willer. The article would not have been possible had Angeline Spain not granted me a couple nights of leave from care of our newborn.

REFERENCES

- Bronfman, A. and Dowell, J. (forthcoming). 'Contextualism About Deontic Conditionals', in N. Charlow and M. Chrisman (eds) *Deontic Modals*. Oxford: Oxford University Press.
- Buchak, L. (2014). *Risk and Rationality*. Oxford: Oxford University Press.
- Cariani, F. (2009). 'The Semantics of "Ought" and the Unity of Modal Discourse', PhD thesis, UC Berkeley.
- Cariani, F. (forthcoming). 'Deontic Modals and Probabilities: One Theory to Rule Them All?', in N. Charlow and M. Chrisman (eds). *Deontic Modals*. Oxford: Oxford University Press.
- Cariani, F., Kaufmann, M. and Kaufmann, S. (2013). 'Deliberative Modality under Epistemic Uncertainty', *Linguistics and Philosophy*, 36, pp. 225–259.
- Carr, J. (2012). 'Deontic Modals without Decision Theory', in *Proceedings of Sinn und Bedeutung* 17, pp. 167–182.
- Charlow, N. (2013). 'What We Know and What To Do', *Synthese* 190, pp. 2291–2323.
- Charlow, N. (forthcoming). 'Decision Theory? Yes! Truth Conditions? No!', in N. Charlow and M. Chrisman (eds) *Deontic Modals*. Oxford: Oxford University Press.

- Dowell, J. (2012). 'Contextualist Solutions to Three Puzzles about Practical Conditionals', in R. Shafer-Landau (ed.) *Oxford Studies in Metaethics*, vol. 7. Oxford: Oxford University Press, pp. 271–303.
- Egan, A., Hawthorne, J. and Weatherson, B. (2005). 'Epistemic Modals in Context', in G. Preyer and P. Peter (eds) *Contextualism in Philosophy*. Oxford: Oxford University Press, pp. 131–170.
- Fara, D. (2003). 'Desires, Scope and Tense', *Philosophical Perspectives* 17(1), pp. 141–163.
- Fara, D. (2013). 'Specifying Desires', *Noûs* 46(2), pp. 250–272.
- von Fintel, K. (unpublished). 'The Best We Can (Expect ot) Get? Challenges to the Classic Semantics for Deontic Modals'. Presented at the 2012 Central APA, Chicago, IL. Available online at: <http://mit.edu/fintel/fintel-2012-apa-ought.pdf>
- Goble, L. (1996). 'Utilitarian Deontic Logic', *Philosophical Studies* 82(3), pp. 317–357.
- Hamblin, C. L. (1959). 'The Modal "Probably"', *Mind* 68(270), pp. 234–240.
- Hintikka, J. (1962). *Knowledge and Belief: An Introduction to the Logic of the Two Notions*. Ithaca, NY: Cornell University Press.
- Horty, J. F. (2001). *Agency and Deontic Logic*. Oxford: Oxford University Press.
- Jackson, F. (1991). 'Decision Theoretic Consequentialism and the Nearest Dearest Objection', *Ethics* 101(3), pp. 461–482.
- Jackson, F. and Pargetter, R. (1986). 'Oughts, Options and Actualism', *The Philosophical Review* 95(2), pp. 233–255.
- Kolodny, N. and MacFarlane, J. (2010). 'Ifs and oughts', *Journal of Philosophy* 107(3), pp. 115–143.
- Kratzer, A. (1977). 'What "Must" and "Can" Must and Can Mean', *Linguistics and Philosophy* 1(3), pp. 337–355.
- Kratzer, A. (1981). 'The Notional Category of Modality', in B. Partee and P. Portner (eds) *Formal Semantics: The Essential Reading*. Oxford: Blackwell.
- Kratzer, A. (2012). *Modals and Conditionals*. Oxford: Oxford University Press.
- Lassiter, D. (2011). *Measurement and Modality: The Scalar Basis of Modal Semantics*, PhD thesis, NYU.
- Lassiter, D. (2014). 'Epistemic: Deontic: Additive: Intermediate', *Pacific Philosophical Quarterly* (this issue).
- Lewis, D. K. (1973). *Counterfactuals*. Oxford: Blackwell.
- Lewis, D. K. (1978). 'Reply to McMichael', *Analysis* 38(2), pp. 85–86.
- MacFarlane, J. (2014). *Assessment Sensitivity*. Oxford: Oxford University Press.
- Ninan, D. (2010). 'Semantics and the Objects of Assertion', *Linguistics and Philosophy* 33(5), pp. 355–380.
- Prior, A. N. (1956). 'The Consequences of Actions', *Proceedings of the Aristotelian Society Supplementary Vol. 30*, pp. 91–99.
- Regan, D. (1980). *Utilitarianism and Cooperation*. Oxford: Oxford University Press.
- Stephenson, T. (2007). 'Judge Dependence, Epistemic Modals, and Predicates of Personal Taste', *Linguistics and Philosophy* 30(4), pp. 487–525.
- Swanson, E. (2010). 'Lessons from the Context Sensitivity of Causal Talk', *Journal of Philosophy* 107(5), pp. 221–242.
- Wedgwood, R. (forthcoming). 'Subjective and Objective Ought', in N. Charlow and M. Chrisman (eds) *Deontic Modals*. Oxford: Oxford University Press.
- Willer, M. (2012). 'A Note on Iffy Oughts', *Journal of Philosophy* 109, pp. 449–461.
- Yalcin, S. (2007). 'Epistemic Modals', *Mind*, 116(4), pp. 983–1027.
- Yalcin, S. (2010). 'Probability Operators', *Philosophy Compass* 5(11), pp. 916–937.
- Yalcin, S. (2012). 'Bayesian Expressivism', *Proceedings of the Aristotelian Society* 112, pp. 123–160.