'Ought': OUT OF ORDER

Stephen Finlay

ABSTRACT: I argue that Angelika Kratzer's introduction of an ordering source parameter into the semantics for natural language modals was a mistake, at least for normative modals. A simpler semantics in a dyadic framework, motivated by the need for a satisfactory treatment of instrumental (or "anankastic") conditionals, also provides the resources for a better accommodation of gradability and "weak necessity", information-sensitivity, and conflicts, with three moves: (i) an end-relational analysis of normative modality, (ii) an analysis of 'ought' or "weak necessity" in terms of *most*, and (iii) appeal to the same pragmatic resources utilized by Kratzer. The paper ends with metasemantic observations about what we should want from a semantics for 'ought'.

The semantics of 'ought' and related modal verbs in natural language is not for the faint of heart. A central desideratum is a *unifying* theory accommodating the many diverse uses, both normative (moral, instrumental, rational, legal, etc.) and nonnormative (logical, metaphysical, nomological, epistemic, dispositional, etc.) This places it at the intersection of many issues in different subfields of philosophy, linguistics, and logic, at least, and so presumably no single scholar could acquire all relevant expertise. Angelika Kratzer's (David Lewis-influenced) *ordering semantics*, which today is widely regarded as orthodoxy, is therefore a remarkable and audacious achievement. Whereas the Lewis-Kratzer semantics was developed from a primary focus on counterfactuals then extended to normative and other uses, this paper offers a metaethicist's perspective, taking as its primary cues the behavior of English modal verbs in normative sentences.

When I first developed my own theory of the meaning of 'ought' in 2005-6 (initially published as "Oughts and Ends" in 2008), like other metaethicists I was ignorant of Kratzer's seminal work on modals in linguistics (1978, 1981, 1991). This defect in metaethics has since been corrected, and

today philosophical work on normative modals typically adopts the Lewis-Kratzer framework. But this new attention has uncovered serious difficulties for its application to normative sentences, stimulating an explosion of work at the intersection of metaethics and linguistics. Whereas these efforts at repairing or replacing the Lewis-Kratzer framework invariably propose additional semantic complexity, I'll argue here that these difficulties are all avoided more straightforwardly and naturally by the simpler semantics I've advanced in "Oughts and Ends" and subsequent work (2010a, 2010b, 2014), a version of the *dyadic* semantics which the Lewis-Kratzer approach supplanted. I'll argue that Kratzer's signature innovation of an *ordering source* parameter, though ingenious, introduced unnecessary complexity into the semantics of modals that we're better off without, at least for English auxiliaries like 'ought', 'must', and 'may'. The simpler semantics has the resources to address the issues motivating ordering semantics, without its weaknesses.

After sketching the central features and motivations of the ordering semantics in Section 1, I'll programmatically investigate problems arising from (i) instrumental conditionals (Section 2)—which motivate my rival approach to normative modality (Section 3), (ii) gradability and "weak necessity" (Section 4), (iii) information-sensitivity (Section 5), and (iv) conflicts (Section 6). I will argue that a simpler dyadic approach addresses these problems at least as well as the ordering semantics, given three moves: (a) an *end-relational* analysis of normative modality, (b) an analysis of 'ought' or "weak necessity" in terms of less-than-universal quantification (*most* rather than *all*), and (c) appeal to the same resources of conversational pragmatics utilized by Kratzer. I conclude (Section 7) by addressing metasemantic objections to my claims about simplicity, with observations about what we should want from a semantics for 'ought'.

1. Ordering Semantics: Dyadic to Polyadic

By orthodoxy and with great success, modal auxiliaries are analyzed as quantifiers over possibilities. Necessity verbs like 'must' and 'have to' are universal quantifiers meaning roughly *in all possibilities*, while possibility verbs like 'may', 'might', 'can', and 'could' are existential quantifiers meaning roughly *in some possibilities*. Different kinds of modality result from defining the relevant domain differently: by consistency with physical laws (nomological), bodies of evidence (epistemic), codes of laws or rules (deontic), etc.

Dyadic semantics classically capture this variability by postulating that in addition to their scope, modal verbs take a second argument for a domain restrictor. Following Kratzer, this *modal base* is identified as a conversational background, usually implicit in the context but articulable with a 'given…' clause. Formally, a modal base is a function *f* from a context *w* to a "premise-set" of propositions *f(w)*. What is *necessary* in *w* given *f* is what logically follows from *f(w)*, and what is *possible* in *w* given *f* is what is logically consistent with *f(w)*. On a simple dyadic approach, this argument-place is undiscriminating. It can take backgrounds that are (i) "realistic" or factual, like *given the circumstances at time t* or *given what subject s knows*, (ii) "nonrealistic"—including normative backgrounds—like *given s's desired ends*, *given what is prescribed by the code of laws L*, or *given what s believes*, or (iii) mixed. But beginning with her (1981), Kratzer insists that "realistic and normative backgrounds need to be kept separate" (2012: 38). She offers two reasons.

First, her "most important argument" is the need to accommodate the *gradability* of modality, as in degrees of deontic ideality, epistemic likelihood, and counterfactual closeness. A metaethically important case is the modal force of 'ought' (or 'should'), which is stronger than that of possibility modals like 'may', but weaker than that of necessity modals like 'must', and is therefore labeled "weak necessity".

Second, normative and other nonrealistic modalities often involve *inconsistency*: ideals can conflict either with each other (e.g. incompatible goals or contradictory laws), or with salient circumstances, (e.g. what one ought to do given violation of some norm, or "contrary-to-duty obligations"). Since anything follows from an inconsistent premise-set, a dyadic semantics predicts that relative to such backgrounds *Must p* is true for any arbitrary *p*, and, since nothing is consistent with an inconsistent set, that *May p* is false for any arbitrary *p*. But inconsistent ideals don't seem to render normative claims trivially true or false.

Kratzer's ingenious solution to both problems is to suggest that modal verbs take a second conversational background, an *ordering source*. This is to move from a dyadic to a *polyadic* semantics (although "dyadic" is often used in the literature, confusingly, where polyadic is meant). This ordering source is also a variable function, *g*, from a context *w* to a premise-set of propositions *g(w)*, comprised of goals, ideals, laws, etc. But rather than restricting the domain, *g* orders it, roughly according to how closely the possibilities approximate the premise-set. The *modal force* of a verb consists in how it quantifies over a subset of possibilities picked from the ordering by a selection

function. Simplifying slightly, *Must p* says, relative to the *f* and *g* in *w*, that *p* is true in *all* possibilities consistent with *f(w)* that are highest-ranked by *g(w)*, while *May p* says that *p* is true in *some* possibilities consistent with *f(w)* that are highest-ranked by *g(w)*. In this paper I use the term *ordering semantics* narrowly to refer only to theories positing an ordering source parameter.

These two kinds of background are utilized in different ways to account for different modal flavors. The modal base is stipulated always to be realistic, on Kratzer's division of labor. Normative or deontic flavors are generated by normative ordering sources, whereas bare alethic (e.g. logical, metaphysical) modals have empty ordering sources, so that all possibilities in the domain are equally ranked. Other modalities (epistemic, counterfactual, etc.) are distinguished by particular combinations of backgrounds. We'll see that normative modalities pose significant difficulties for this shift from dyadic semantics to polyadic ordering semantics.

2. Instrumental Conditionals

Consider *instrumental conditionals*,[1] like

> (1) If you want to go to Harlem, you must take the A train.
> (2) If you are (going) to go to Harlem, you must take the A train.
> (3) (In order) to go to Harlem, you must take the A train.

(I begin with 'must', as 'ought' introduces additional complications.) Surface differences notwithstanding, these are widely agreed to have instrumental readings as saying, approximately, that taking the A train is a necessary means to going to Harlem, and to have the normative flavor of a teleological modal, with the force of hypothetical advice.

The Lewis-Kratzer semantics has difficulties accounting for these conditionals, as first observed by Kjell Johan Sæbø (2001). This might seem a minor nuisance posed by an obscure subclass of sentences; Sæbø suggests they have "attracted little attention in linguistics and philosophy" (2001: 428). I believe this impression is mistaken, for the reason that instrumental conditionals are key to

---

[1] (Some of) these are commonly called "anankastic conditionals" by linguists and "hypothetical imperatives" by philosophers; both labels are misnomers. (a) "Anankastic" signifies *necessity/compulsion*, but necessity is neither necessary (e.g. "existential anankastics" like 'If you want to go to Harlem, you could take the A train') nor sufficient (e.g. 'If you are in Harlem, then you must be in New York'). (b) These sentences have declarative rather than imperative mood.

analyzing normative modality quite generally. On a controversial though perennially popular view (e.g. Foot 1972), *every* normative or deontic 'ought' ('must', etc.) sentence is at least implicitly instrumental in some way, even moral *oughts*. If this view is correct, then a satisfactory treatment of instrumental conditionals is essential for a successful semantics for normative modals. I'll first sketch the problems these sentences pose for the ordering semantics, and then briefly explain how I've argued, in previous work, that these can be naturally resolved by returning to a dyadic semantics. I won't attempt here to establish the success of this proposal; interested or skeptical readers may examine my earlier work. Rather, I shall make the case that this simpler semantics motivated by the need for a satisfactory treatment of instrumental conditionals also compares favorably against the ordering semantics with respect to the further features of gradability, conflicts, and information-sensitivity—despite (two of) these being Kratzer's primary motivations for adopting ordering semantics over a dyadic framework.

Focus first on 'If *s* wants…' sentences like (1). By Kratzer's own influential "restrictor" semantics for conditionals, 'if p' functions to update the modal base f with the antecedent p, yielding the restricted domain $f+p(w)$. This suggests reading (1) as saying, for salient f and g, that *in all possibilities consistent with f in which you want to go to Harlem that are highest-ranked by g, you take the A train.* However, there's no guarantee that the salient ordering source g will rank go-to-Harlem worlds highest, so restricting to worlds where you have this desire doesn't entail that in all the best remaining worlds you take the A train, even if it's the only way to get to Harlem. Intuitively, the goal *you go to Harlem* should instead be determining which possibilities are best. This presents the ordering semantics with (at least) two puzzles:

(i)   How can the antecedent clause influence the ordering source?

(ii)   How does the reference to desire or purpose contribute to sentences like (1)—since the relevant goal is the "internal antecedent", *(you) go to Harlem*?

For reasons of space I won't focus on the second puzzle here. There are ample grounds for thinking that the reference to desires doesn't make a straightforwardly compositional contribution to truth conditions,[2] as most writers have concluded and I will here assume.[3]

Sæbø responds to the first puzzle by proposing, radically, that 'if' is ambiguous, and can update either the modal base or the ordering source. On its instrumental reading, the antecedent of (1) would update the background ordering source g with the internal antecedent, that e: you go to Harlem. This fix isn't sufficient, however. Merely *updating* an ordering source with the goal e doesn't guarantee that all the $g+e$-best possibilities will be go-to-Harlem worlds, because of a problem from *inconsistent goals*: e could be incompatible with other goals in the salient g, like going to Hoboken (perhaps you have a grandmother who lives in each place).[4] A sentence like (4) could then be true in such circumstances:

(4) #To go to Harlem, you may go to Hoboken instead.

The "designated goal" in the antecedent clearly needs to take priority over any other goals or ideals. One suggestion is that it interacts with the preliminary ordering source, g_0, by eliminating anything inconsistent with it, thereby guaranteeing that the $g+e$-best possibilities will all be go-to-Harlem worlds.[5] However, this encounters a problem from *consistent goals*. Suppose you happen to have the desire or goal to kiss the footballer Ruud van Nistelrooy, who will be on the A train. Then it's predicted that (1) may be true even if there are other, equally reliable ways to go to Harlem (Huitink 2005)—and worse, that (5) may be true (Nissenbaum 2005);

(5) #To go to Harlem, you have to kiss Ruud van Nistelrooy.

---

[2] First, the contrapositive of (1) is apparently 'If you don't take the A train, you can't go to Harlem', rather than 'If you don't take the A train, you can't want to go to Harlem' (Sæbø 2001: 427).

Second, (1) seems roughly equivalent to (2) and (3). While these other constructions can sometimes indicate an agent's purpose, this doesn't seem to be their essential function here (Finlay 2010b). Nissenbaum 2005 proposes the gloss: *you must take the A train with the purpose of going to Harlem*; von Fintel & Iatridou 2005 offer convincing objections.

Third, these sentences exhibit anomalous *detaching* behavior, resisting modus ponens. One can coherently assert 'If Henry wants to be a famous mass murderer, then he has to kill a lot of people', while refusing to assert 'Henry has to kill a lot of people' upon learning that Henry has this disturbing desire.

[3] Condoravdi & Lauer (forthcoming) offer a dissenting view, turning on a reading of 'want' that I would argue is artificial.

[4] von Fintel & Iatridou 2005, von Stechow et al 2005.

[5] Proposed in an early version of von Fintel & Iatridou 2005.

To solve these two problems of inconsistent and consistent goals, Janneke Huitink (2005) proposes that the antecedent functions to *constitute* the ordering source, rather than merely update it. Since ordering sources tolerate inconsistencies by design, and conditionals normally function to update backgrounds, this proposal seems ad hoc and anomalous.[6] It also throws the baby out with the bathwater, since instrumental conditionals do sometimes exhibit sensitivity to other goals or ideals, at least in the case of 'ought' variants like (6), though arguably not for 'must'.[7]

(6) To go to Harlem, you ought to take the A train. It's cheapest.

The favored solution to these problems, proposed by Kai von Fintel and Sabine Iatridou (2008), introduces further semantic complexity by adopting Kratzer's idea of multiple ordering sources. They suggest that 'ought' differs from 'must' in taking a secondary ordering source, having the form $ought^{f,g_1,g_2}(p)$ with roughly the modal force: *in all possibilities consistent with f highest-ranked by g_1 that are highest-ranked by g_2*. For instrumental conditionals, the designated goal determines the primary ordering source g_1, with other salient goals or ideals providing the secondary ordering source g_2. (6) therefore says that in all possibilities highest-ranked with regard to going to Harlem that are also highest-ranked by the secondary ideals (e.g. cost, comfort), you take the A train. This analysis solves many problems, though still predicts the truth of (7) in the van Nistelrooy scenario, which they suggest—implausibly, I think—is infelicitous rather than false:

(7) #To go to Harlem, you ought to kiss Ruud van Nistelrooy.

However, a problem of *unattainable goals* remains for any analysis that assigns the designated goal to an ordering source.[8] By design, ordering sources needn't be consistent with modal bases. Suppose there are no possibilities consistent with *f* in which you go to Harlem; e.g. you're locked up in jail. The singleton set of ideals {you go to Harlem} still generates an ordering: all possibilities in the modal base will be equally best. So the ordering semantics seems to predict incorrectly that the following sentences would be true in this context:

---

[6] As von Stechow et al. 2005 point out, it is also (almost) functionally equivalent to eliminating ordering sources altogether, as I propose below.

[7] This is contested; see discussion of the "Chinese train" example in von Stechow et al. 2005. The correct response, I believe, appeals to ellipsis in the antecedent (von Fintel & Iatridou 2005: 16). The existence of inferior means justifies challenges to and retractions of 'must' claims; e.g. A: 'No, you don't *have* to take the Chinese train to go to Vladivostok; you *could* take the Russian train.'; B: 'Yes, but you do have to take the Chinese train to go to Vladivostok *comfortably*.'

[8] Cf. von Stechow et al. 2005: 8.

(8) #To go to Harlem, you must not go to Harlem.

(9) #To go to Harlem, you must stay in jail.

Observe that these sentences seem false rather than merely infelicitous, unlike their purported translation,

(10) In all accessible possibilities which are closest to the ideal that you to go to Harlem, you don't go to Harlem/ you stay in jail.

My claim here is not that these problems are fatal for ordering semantics. One might for example postulate a (controversial but popular) "diversity condition" in the semantics,[9] requiring that the scope of a modal operator be neither entailed nor ruled out by *f* whenever *g* isn't empty. But I wish to point out how much more simply and naturally a dyadic semantic analysis of instrumental conditionals avoids these problems.

Notice that the challenge confronting ordering semantics here is, effectively, to explain how a goal in an ordering source could behave as if it were in the modal base instead. The truth conditions of instrumental conditionals, like (1)-(3), seem sensitive only to possibilities consistent with achieving the designated goal. So suppose we simply allow the conditional to function in its ordinary way, unambiguously adding the antecedent's designated goal to the modal base instead.[10] I have argued (especially Finlay 2010b, 2014) that a fully unifying, compositional, and conservative analysis then comes into view at least for sentences like (2) with 'if *s* is (going) to…' antecedents, which also supports promising analyses of other grammatical variants like (1) and (3).[11] Briefly:

The prospective aspect in these antecedent clauses encourages a transparently and compositionally *temporal* analysis: to represent an event *e* as "[going] to" happen is to project it into the relative future. A sentence like (2) therefore grammatically represents the modality in the consequent clause as being (temporally or metaphysically) prior to the event *e* in the antecedent; i.e. the *necessity* of taking the A train is located prior to your going to Harlem. This reverses the order we

---

[9] E.g. Condoravdi 2002.

[10] Finlay 2009; see also discussion in von Stechow et al. 2005: 6.

[11] 'In order to' is plausibly an idiomatic device for representing the same reversed temporal relationship described below. I argue for an analysis of 'if *s* wants…' variants as relevance/instrumental conditional hybrids (Finlay 2010b, 2014: 70-1), which explains the anomalies in note 2. Relevance analyses are also tentatively suggested in von Fintel & Iatridou 2005, von Stechow et al. 2005.

read into conditionals by default;[12] for example, the non-instrumental conditional (11) is naturally interpreted as representing the necessity of taking the A train as (temporally) consequent upon going to Harlem:

>   (11) If you go to Harlem, then you must take the A train.

Whereas Kratzer doesn't herself analyze how modality interacts with temporality, I maintain that by virtue of the temporal structure of (1)-(3), 'must' there expresses the necessity of (temporal or metaphysical) *pre-conditions* for the antecedent's event obtaining, and thereby an instrumental (or, roughly, means-end) relationship.

This simple analysis avoids the ordering semantics' problem from inconsistent goals (since it requires consistency with the designated goal) and also the problem from consistent goals (since e.g. kissing van Nistelrooy isn't a necessary condition for going to Harlem), while avoiding any ambiguity in the semantics for 'if'. Of course, we also need an account of the contrast between 'must' and 'ought', and of inconsistencies (e.g. to address the problem of unattainable goals); here I beg the reader's patience until Sections 4 and 6, respectively. But if I am correct, then instrumental conditionals are best analyzed without appeal to any ordering source.

3. The Source of Normativity

Can't the champion of ordering semantics just concede this result, and say that at least some instrumental conditionals involve an empty ordering source? The problem is that instrumental conditionals are almost universally agreed to have an at least minimally normative character—which is largely why moral philosophers puzzle over "hypothetical imperatives" and Kratzerian analyses have almost universally assumed that the designated goal must go into an ordering source. So this concession would apparently acknowledge that at least some normative modals are best analyzed without appeal to ordering sources. I'll now explain why this poses a serious challenge to the ordering semantics.

---

[12] Sæbø 2001 analyzes this as placing the consequent's *event* (the means) prior to the antecedent's event (the end), which fails to distinguish instrumental conditionals from many epistemic conditionals; e.g 'If the ground is wet, then it has to have been raining.' In Finlay 2010b I endorsed these as jointly necessary and sufficient conditions for an instrumental conditional; in Finlay 2014 I rather reject Sæbø's condition.

Kratzer intends her semantics to be neutral toward the metaethical question of what makes an occurrence of a modal verb normative. Answering that question, she says (p.c.), is the philosopher's job. But the ordering semantics is metaethically committal nonetheless, in a potentially problematic way. The stipulation that modal bases are realistic imposes a condition that normativity must enter through an ordering source. Although sometimes writing as if nonempty ordering sources are also sufficient for normative modality (e.g. distinguishing ordering sources as "normative backgrounds" from modal bases as "factual backgrounds"), Kratzer also posits nonnormative ordering sources at least in analyzing some epistemic, dispositional, and counterfactual modals. Rather, a normative flavor supposedly arises from ordering sources with normative content or selected in a normative way.

The dyadic analysis of instrumental conditionals supports a conflicting view at least of the source of an *instrumentally* normative flavor: that it reduces to modality restricted by some state of affairs in the relative future, or what I call an *end*,[13] as included in the modal base. A reasonable hypothesis is that our sense for normative flavor is pragmatically triggered by a sentence's suitability for practical purposes such as giving advice or guiding deliberation: uses for which these end-relational or teleological sentences are especially apt whenever the end is relevantly valued or desired. What could be more apt for purposes of advice than telling someone what is necessary for achieving their desired or intended goal?

This leaves open the possibility that Kratzer's approach is correct for non-instrumentally normative modality, such as the deontic modalities of morality, law, and all-things-considered practical reason. Deontic claims certainly differ in important ways from ordinary "hypothetical imperatives". But plausibly, normativity is one genus of which these are different species. My work in metaethics argues that normative modality just is *end-relational* modality, concerning what is necessary, possible, etc., conditional on the obtaining of some end.[14] This would vindicate the classic view in metaethics that the "categorical imperatives" of morality are really just special kinds of "hypothetical imperatives". If this end-relational theory is correct, then the dyadic analysis of

---

[13] Paul Portner labels these 'priorities'; I prefer 'ends' both for continuity with ancient philosophical tradition, and because 'priority' suggests rather the importance or weight of an end/norm (e.g. as used in Horty 2012).

[14] I analyze the prescriptive, "categorical" character of moral utterances as rhetorical, arising from suppressing reference to the end made salient by the speaker's rather than the agent's desires; see especially Finlay 2014: Ch. 7, and cf. Harman 1996.

instrumental conditionals suggests that all normative modalities, even the moral/deontic, might be best analyzed without ordering sources.[15][16]

Semanticists may be understandably reluctant to commit to controversial philosophical positions, but metaethical neutrality is here an unattainable luxury. The Lewis-Kratzer semantics appears incompatible not only with my end-relational theory, but also with popular rival metaethical views. Contrary to analysis of the moral 'ought' as relative to ordering sources "represent[ing] the content of a body of laws or regulations" (Kratzer 2012: 37), some philosophers argue that the content of moral laws can't even be described or conceived except as laws about what we *ought* to do, where this 'ought' expresses a primitive or irreducible normative relation (e.g. Boghossian 2006: 24f). If this is correct, then attempts to accommodate the moral 'ought' with a unifying relational semantics will lead to vicious circularity: what we morally ought to do is whatever is highest-ranked by the laws concerning what we morally ought to do.

Metaethicists of this stripe also view the extent of *moral disagreement* as refuting claims that moral 'ought's are relativized to particular laws or ideals. Speakers subscribing to different moral laws seem genuinely to disagree when making moral 'ought' claims, which most metaethicists believe wouldn't be the case if each were just speaking about what is highest-ranked by their respectively favored laws.[17] Other metaethicists deny that normative claims quantify over possibilities or express propositions at all. Despite the aspiration to neutrality, the ordering semantics may therefore be compatible with only a narrow range of metaethical theories. Since taking sides in these debates is unavoidable in pursuing a unified modal semantics, we should boldly go wherever the linguistic evidence leads—which, I'll now argue in earnest, is consistently in the direction of a dyadic semantics and the end-relational theory.

---

[15] Why is this theory *dyadic*, if we can distinguish between the end *e*, and the "realistic part" of the modal base, *f-e*? Answer: the modal force of 'ought' (i.e. the operation it performs on its inputs) doesn't differentiate between *e* and anything else in *f*. The end therefore needn't be identified as a separate parameter; see further discussion in Section 7.

[16] Two further supporting considerations: (1) the distinctive deontic logic shared by teleological and deontic modals reduces to ordinary modal logic given only the assumption of conditionalization on a hypothetical outcome (the "Kanger-Anderson" reduction). (2) Teleological and deontic uses of modal verbs share grammatical features of tense which distinguish them from epistemic and other uses. I argue (2010b, 2014: Ch. 3) that the end-relational analysis explains these compositionally, and suggest that in deontic uses they function as grammatical indicators of suppressed conditionalization on an end, which triggers our sense for normativity.

[17] In unpublished work, Janice Dowell proposes analyzing moral 'ought' with ordering sources invoking another normative concept, like: in view of the *most important* standards. This solution depends on these further concepts not being themselves relativistic, but I argue (2014: 252-3) for a parallel treatment. I address the problem of disagreement in Björnsson & Finlay 2010, Finlay 2014: Ch. 8.

4. Gradability and "Weak Necessity"

Kratzer identifies the need to accommodate *gradable* modality, or degrees of possibility (goodness, likelihood, typicality, closeness) as her primary argument for ordering sources. I think it is at least unclear whether such a need exists for English modal auxiliaries like 'must', 'may', and 'ought', which don't naturally take degree modifiers like '-er', '-est', 'most', 'slightly', 'very', etc. (Portner 2009).[18] But there may be a special case for a gradable semantics for 'ought' and 'should', which in both epistemic and normative uses lie somewhere between 'must' and 'may' in logical strength. *Must p* entails *Ought p*, while *Ought p* entails *May p*, and not vice versa. I'll now investigate whether ordering sources are needed to account for this distinctive "weak necessity" of 'ought'.

From the outset, 'ought' also presents a prima facie problem for ordering semantics. Intuitively, 'ought' is to 'must' as 'best' is to 'only' (Sloman 1970), and so *Ought p* is commonly glossed as: *it's best that p*. You *must* take the A train to go to Harlem if it is the only available way, and *ought* to take the A train to go to Harlem if it is the best available way. But because the ordering semantics analyzes 'must' as meaning roughly *in all the best possibilities*, finding logical space for 'ought' presents a difficulty.

Whereas Kratzer doesn't ever propose an analysis of 'ought' herself,[19] various proposals have recently been advanced on behalf of ordering semantics. These all involve introducing additional semantic complexity—such as multiple ordering sources (von Fintel & Iatridou 2008), "merging" operations on these (Charlow 2013, Katz, Portner & Rubinstein 2012), counterfactual conditions (Silk ms.), additional contrast-set parameters (Cariani, Kaufmann & Kaufmann 2013, Charlow 2013, Carr ms.) and special rules of use (Rubinstein 2013). I believe these proposals all face serious obstacles, though I can't demonstrate that here.[20] Instead, I'll show how the dyadic framework provides a natural, simpler way of accommodating weak necessity and degrees of possibility.

---

[18] Kratzer responds that this is a "language-specific fact" and that auxiliaries in some other languages and some English modal adjectives do take such modifiers (2012: 42)—which I don't think justifies attributing a gradable semantics to *English auxiliaries*. Admittedly, there are some ways to qualify these with degree modifiers (see Portner & Rubinstein, this volume), but these seem to apply indiscriminately, even to uncontroversially nongradable terms like 'dead' and 'guitar' (e.g. 'X is more a guitar than Y is.')

[19] P.c. She does offer an analysis of "weak necessity" (1991: 644).

[20] For problems for multiple ordering sources, see Rubinstein 2013. Silk faces the problem that 'ought' isn't synonymous with 'would have to', and a nonarbitrary kind of counterfactual condition is yet to be identified. Rubinstein's suggestion that 'ought' indicates a controversial ordering source conflicts with the Theseus/miner scenarios

Given the equivalence,

'must' : 'ought' : 'may' :: *all* : *???* : *some*,

the naïve and obvious hypothesis is that the modal force of 'ought' (and 'should') is roughly that of *most*, which has appropriately intermediate logical strength.[21] (This suggestion is precisified below). This applies intuitively to the epistemic 'ought'; e.g. that

(12) It ought to rain today,

says roughly that it rains today in most possibilities consistent with the background evidence. On the classical theory of probability as a proportion or measure of possibility-space, (12) would then be equivalent to saying that it *probably* rains today given the evidence, as many writers have accepted.[22] "Weak necessity" would therefore turn out to be *greatest probability*.[23]

This simple analysis is usually raised only to be summarily dismissed. The primary objection offered is the classical theory's difficulties with infinity (e.g. Portner 2009: 32-3). On the Kripkean possible world semantics in which the Lewis-Kratzer theory is modeled, there are infinitely many possible worlds in any typical domain, so the notion of a *count* or *proportion* of possibilities is undefined. But this is merely a choice of formal model, which as Kratzer insists (2012: 10), shouldn't itself dictate the content of semantic theory. Alternative models quantify over possibilities of a more coarsely individuated nature, corresponding to *partitions* on sets of possible worlds, which are more plausibly countable or measurable.[24] Note also that the ordering semantics faces its own problems with infinity. To avoid the artificial Limit Assumption that there always are *best* worlds,

---

below, in which relevant ideals are uncontroversial, and with moral uses of 'must' which are often controversial. See Section 6 for some objections to other accounts.

[21] This is called a "common intuition" in Portner 2009: 32, and the "traditional view" in Copley 2006: 4, but I'm unaware of anyone else defending it in print. Copley attributes it to Larry Horn, who reports (p.c.) maintaining the view outside of print. Previously I emphasized *probably*, tentatively analyzed in terms of *most* (2010a: 80; 2014: 73); see also Wheeler 1974, 2013.

[22] Some common objections: (i) *Ought p but not p* is acceptable, unlike *Probably p but not p* (e.g. Copley 2006). Reply: unlike tenseless modal adverbs, auxiliaries like 'ought' needn't be relativized to full present evidence (Thomson 2008, Finlay 2014: 73, Wedgwood, this volume). (ii) *Ought p* seems bad when *p* is only incrementally more likely than some relevant alternatives. Reply: it may just be infelicitous; compare *Most likely, p.*

[23] Cf. Kratzer 1991: 644.

[24] Finlay 2014: 44. (Some writers, like Cariani, Kaufmann & Kaufmann 2013, Dowell 2013, Charlow 2013, induce such partitions in a possible worlds framework, to solve problems for ordering semantics; see below). Another option is to adopt a measure function defined for infinity.

Kratzer advances definitions for the simplest auxiliaries like 'must' and 'may' that she herself admits are "more complicated than might seem necessary" (2012: 40).[25]

Reluctance to commit to a controversial theory of probability is understandable, but this may be another place where modal semantics can't afford the luxury of neutrality. Observing that the semantics for 'ought' needs to be sensitive to probabilities, multiple writers propose replacing the modal base with an argument-place for a quantitative probability distribution over possible worlds, usually also replacing the ordering source with a value function.[26] On the classical theory such a parameter is unnecessary, as a modal base is sufficient to determine a unique probability function.[27] I myself am optimistic that this approach can be vindicated,[28] but how can semantic theorizing proceed when faced with such difficulties?

I suggest we bear in mind that our subject is the meaning of natural language. Whatever philosophical challenges it faces, the classical theory of probability has excellent claim to be the naïve theory of common sense, which is good reason not to build anything more sophisticated into our semantics, "beyond the limits of what the faculty of language provides for everyone" (Kratzer 2012: 43).[29] I find it implausible that ordinary modal speech and thought quantifies over maximally fine-grained possible *worlds*. Given only the information that a group of miners is trapped in one of two shafts, for example, it's extremely natural to say that there are *two* relevant possibilities, and that relative to this information each therefore has the probability ½. (It also seems natural to represent probabilities geometrically rather than arithmetically, treating *most* as a measure of possibility-space

---

[25] Formally:

$[[\text{must } p]]^{f,g} = \text{T}$ iff for all $u \in \cap f(w)$ there is a $v \in \cap f(w)$ such that $v \leq_{g(w)} u$, and for all $z \in \cap f(w)$, if $z \leq_{g(w)} v$, then $z \rightarrow p$.

A formalization of a basic multiple ordering source semantics for 'ought' that avoids the Limit Assumption:

$[[\text{ought } p]]^{f,g1,g2} = \text{T}$ iff for all $u \in \cap f(w)$ there is a $v \in \cap f(w)$ such that $v \leq_{g1(w)} u$, and for all $z \in \cap f(w)$, if $z \leq_{g1(w)} v$, then there is a $q \in \cap f(w)$ such that $q \leq_{g1(w)} z$ and $q \leq_{g2(w)} z$, and for all $r \in \cap f(w)$, if $r \leq_{g1(w)} q$ and $r \leq_{g2(w)} q$, then $r \rightarrow p$.

Swanson 2011 presents a scenario that requires even further complications to the ordering semantics.

[26] For example, Goble 1996, Wedgwood 2006, this volume, Cariani this volume.

[27] Cf. Williamson 2000, Kolodny & MacFarlane 2010. Probabilistic information can be included in modal bases where necessary.

[28] For an amateur attempt, see Finlay 2014: 44-5. One constraint on an acceptable model is that relevant possibilities needn't be known.

[29] This claim is defended in Section 7. Ordering semantics are often defended against quantitative frameworks on similar grounds (e.g. Kratzer 2012: 25).

rather than a count of discrete possibilities). At the least, I suggest that the theoretical virtues demonstrated below for the *most* analysis warrant a sympathetic reconsideration of these difficulties.

A second objection is that this simple hypothesis can't be extended to the *normative* uses of 'ought': what a person ought to do isn't what they probably do (e.g. Portner 2009: 76-7). But in "Oughts and Ends" I argue that it actually yields a promising analysis of normative 'ought's if we adopt a *contrastivist* semantics. This interprets 'ought' claims relative to a contrast set R of mutually incompatible propositions, or *options*, $\{r_1, r_2, …, r_n\}$, which we can treat as generated by a conversational background o. (While this adds a second parameter, it is specific to 'ought', for which ordering semantics generally introduce third or more parameters.[30] For this reason, I think it is reasonable to continue describing this as a move within the dyadic semantics framework.) For the practical 'ought' used in advice and deliberation, this background o is articulated with a phrase like 'rather than any other option in the agent's power to perform at t'. My proposal is that 'ought' has the modal force of comparative probability, meaning *in more possibilities* (or: *more likely*) *than any alternative, r_i, in $o(w)$*.[31]

To see the application to normative modalities, consider a teleological 'ought' like

(13) To evade arrest, Max ought to mingle with the crowd.

Following the blueprint of the end-relational analysis of 'must' above, we assume a preliminary modal base f of circumstances or evidence, which is then updated with the end e: that Max subsequently evades arrest. As with the teleological 'must' and 'may' (and also ability modals, following Kratzer [1981: 53]), we exclude from f any facts or information about the agent's psychological dispositions to choose any one option over any other. Suppose Max is disposed not to mingle with the crowd because he falsely believes it wouldn't succeed; we wouldn't on those grounds dissent with the assertion of an instrumental possibility sentence like (14):

---

[30] We could dispense with this parameter by treating $o(w)$ as fixed by p plus context, but a contrast parameter for 'ought' is independently well-motivated, embraced for numerous reasons; see Sloman 1970, Jackson 1985, Cariani 2013a, 2013b, Snedegar 2012, Finlay & Snedegar 2014.

[31] My simple dyadic semantics can be formally expressed as:

$[[\text{must } p]]^f$ = T iff for ALL $u \in \cap f(w)$: $u \to p$.

$[[\text{ought } p]]^{f,o}$ = T iff for MOST$_o$ $u \in \cap f(w)$: $u \to p$; or:

$[[\text{ought } p]]^{f,o}$ = T iff for all $r \in o(w)$: $|\{u \in \cap f(w): u \to p\}| > |\{v \in \cap f(w): v \to r\}|$.

(Compare note 25). Whereas I advanced this analysis of 'ought' primarily to get the right truth-conditions for normative uses, note that Yalcin (2010) has since advanced a very similar comparative analysis for 'probably'.

(14) To evade arrest, Max could mingle with the crowd.

Relative to a preliminary modal base of this kind, every option is equally represented in the possibility-space; I call this *symmetry of choice*.[32]

Updating this preliminary background *f*, characterized by symmetry of choice, with the end *e* yields the following result: the option chosen in most of the remaining possibility-space *f+e(w)* is the option on the choice of which *e* eventuates in more of the possibility-space than it does on the choice of any alternative r_i. In simpler language, the option most likely chosen is that which, if chosen, would make the end most likely. Formally, given symmetry of choice, *Equivalence* holds:

Equivalence: $\forall r_i \in R$: pr($p|f+e$) > pr($r_i|f+e$) iff pr($e|f+p$) > pr($e|f+r_i$).

Equivalence might seem too complex to be playing any role in ordinary normative thought. But represented geometrically it's highly intuitive:
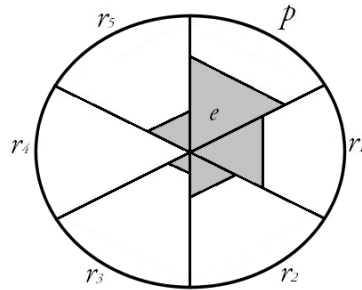
As is obvious at a glance, once we restrict to the *e*-space (shaded), more of the remainder is *p*-space than r_i-space, for any *i*. Informally, this analysis says that whereas what *must* be done in order to achieve some end is whatever is necessary for the end, what *ought* to be done in order to achieve the end is just whatever is the *surest* or *most reliable* way to achieve it.[33]

---

[32] A naturalized version of Kant's assumption of transcendental freedom: to deliberate, one must regard oneself as equally free to choose any available option.

[33] Cariani (2013b: 76) offers two objections to symmetry of choice. First, "it is plausible to assume that the contextually supplied probability function might be either some salient credence…or an evidential probability function." Reply: this is to reject without argument my claim (and Kant's) that thought about whether S ought to do A essentially ignores any evidence that S will do A, by virtue of its very nature as *normative*. Second, "there is no guarantee that a set of alternatives equiprobable relative to an initial background will remain equiprobable after the background is updated," e.g. by a conditional like *if it's snowing outside*. Reply: symmetry of choice is always (re-)applied *after* any circumstantial updating, for the reason that the assumption of transcendental freedom applies precisely in the circumstance of choice. Another objection to the *most* analysis, pressed by Cariani and Ralph Wedgwood (p.c.), is that it entails failure of Agglomeration: *Op* & *Oq* → *O(p* & *q)*. I believe this may rather be another virtue; see Jackson 1985.

Promisingly, this captures the difference between *must* and *ought* in many ordinary scenarios. But can it be extended to an analysis of all normative 'ought' sentences? An obvious objection is that what ought to be done, even instrumentally, isn't always just what makes some primary end more likely, but frequently depends on multiple, sometimes competing ends or desiderata, like comfort, decency, cost, or safety. Here we confront Kratzer's other reason for introducing ordering sources. While I won't address the issue of multiple ends or conflicts until Section 6, I'll now show that this semantics at least offers a superior treatment of simple cases, involving single ends.

Consider the following scenario:[34]

> *Single-Minded Theseus*: Theseus is rushing through the Labyrinth with a single goal in mind: to find and kill the minotaur before it slaughters his compatriots. He reaches a room with three doors, with the following probabilities of success: door A=.5, door B=.3, door C=0.

My dyadic semantics generates the following, intuitive verdict:

> (15) To save his compatriots, Theseus *ought* to choose A, though he doesn't *have to*, as he *could* choose B, but he *must* choose one of A or B, and *mustn't* choose C.

The ordering semantics is unable to yield these results so simply or naturally. Perhaps the most promising strategy (von Fintel & Iatridou 2005) is to treat reliability as a secondary ideal for the 'ought' claim, utilizing multiple ordering sources. But reliability isn't normatively optional in the way secondary ideals are. Even if the *only* thing that matters is saving his compatriots, Theseus ought to take door A. Reliability might be thought a secondary ideal because it can be traded off against cost or comfort, etc., but this is really to trade off the importance of *success*, i.e. of the end itself, against conflicting ideals.

This end-relational semantics has an additional advantage. It accommodates degrees of possibility through the (classical) probabilistic structure of possibility-spaces, which in effect directly induces an ordering of the *options* in the contrast set. Importantly, this differs from the gradability provided by ordering sources in the Lewis-Kratzer semantics, which instead directly induce an order on *possible worlds* according to their approximation to the ideal. So while the *most* analysis simply identifies the "best" option as the most likely/reliable option, the ordering semantics rather

---

[34] Finlay 2010a: 81n, 2014: 77.

identifies an option as "best" just in case (roughly) it is true in all the best worlds. This indirectness is problematic: what is true in the best worlds might also be true in the very worst worlds, and it seems arbitrary to determine an option's value by looking exclusively at best worlds and ignoring worst worlds.[35] As we'll see below, some recent work in ordering semantics therefore seeks to induce orderings on *sets* of possible worlds, partitioned by options—by introducing a contrast-set parameter (e.g. Cariani 2013a), a special selection function (e.g. Charlow 2013), or a special kind of ordering source (e.g. Dowell 2013). The *most* analysis delivers these desired results for free.

5. Information-Sensitivity

Further difficulties for the ordering semantics are observed in how normative modals are sometimes sensitive to subjects' information. Whereas Kratzer originally analyzes normative modals as always taking *circumstantial* modal bases, picking out premise-sets of facts objectively, some normative claims are subjective or information-relative. A simple amendment is to allow modalities with *epistemic* modal bases, picking out sets of facts that constitute what the relevant subject knows ("information-states").[36] But this may not be enough, as there are also more subjective normative uses of modal verbs, such as the 'ought' of practical reason, which is sensitive to what is merely *believed* rather than known.

Kratzer proposes (p.c.) to accommodate these subjective uses with a circumstantial modal base of *facts about* the agent's beliefs, and an ordering source consisting of norms of rationality, which I'll call a *subjectivized end* strategy. While I agree that we sometimes make claims about what an agent ought rationally to do given his beliefs, this fits uncomfortably with other ordinary subjective 'ought' claims. First, we often refer to ordinary, objective ends when making subjective 'ought' claims. Suppose for example that Theseus's beliefs favor choosing door A, although this option is in fact incompatible with achieving his goal. We might then say,

(16) Given what he believes, *to save his compatriots* Theseus ought to take door A.

Second, it makes wrong predictions about appropriate modal force. There is plausibly no separation between what is rationally necessary and what is rationally optimal. In order to act

---

[35] Cf. Jackson 1985. This selection function is also what forces the Limit Assumption, which the *most* analysis avoids without fuss.

[36] Finlay 2009: 322n, Kolodny & MacFarlane 2010: 130.

rationally, Theseus *must* take door A. It seems infelicitously weak to say that to act rationally he *ought* to take door A, and simply false to say that to act rationally he *could* take door B. By contrast it's natural to say that given what he believes, to save his compatriots Theseus *ought* to take door A but *doesn't have to*, and *could* take door B though he *oughtn't*.

Once Kratzer's realism constraint is abandoned, we can simply appeal to a modal base consisting of the (nonrealistic) contents of the subject's beliefs ("given *what* you believe") to account for these subjective normative claims.[37] The realism constraint seems gratuitous, in any case. While the modal base has the function of domain restrictor in the semantics, sometimes conversationally relevant modalities are restricted rather than merely ordered by nonfactual information.[38] Instead of appealing to norms of rationality, the dyadic semantics allows us simply to appeal to the same objective ends (e.g. *to save his compatriots*), and account for the subjective 'must', 'may', and 'ought' as concerning what is respectively necessary, possible, and most likely relative to the subject's information i and the end e. Consider a subjectivized version of *Single-Minded Theseus*, where the probabilities are subjective (i.e. relative to Theseus's beliefs); my semantics then predicts the following, intuitive verdict, parallel to (15):

> (17) Given what he believes, to save his compatriots Theseus *ought* to choose A, though he doesn't *have to*, as he *could* choose B (albeit foolishly!), but he *must* choose either A or B, and *mustn't* choose C.

Unlike Kratzer's proposed solution, this semantics generates the appropriate judgments and respects the relativization to objective ends. It requires no appeal to controversial norms of rationality (indeed, I think it might rather explain them), and directly explains why what one ought to do *given complete accurate beliefs* converges on what one ought to do *given the facts*: these different functions pick out identical premise-sets.

---

[37] Finlay 2010a: 84. With inconsistent beliefs the modal base must be restricted to a consistent subset (Section 6).

[38] E.g. exocentric epistemic claims. Consider a variation of von Fintel & Gillies' mastermind scenario where a player accepts false information which she realizes entail there are two reds. A nondeceived observer can correctly say, 'She knows there must be two reds'; i.e. *must given her (false) information*.

A second, more widely discussed problem about information-sensitivity involves the way that strengthening information can change what an agent ought to do, by changing which options are best. Niko Kolodny and John MacFarlane (2010) illustrate this with the following scenario:[39]

> *Miners:* Ten miners are trapped together in one of two shafts, A and B, in danger of drowning in impending floodwaters. We can block either shaft, but not both, which will save all the miners if they're in the blocked shaft but drown them otherwise. Or we can block neither shaft, which is certain to save nine miners and drown one.

In a context of this information, the following sentences are all typically judged true:

(18a) If the miners are all in shaft A, then we ought to block A.

(18b) If the miners are all in shaft B, then we ought to block B.

(18c) Either the miners are all in A or they are all in B.

(18d) We ought to block neither shaft.

This presents a puzzle. The conditional clauses of (18a) and (18b) apparently just strengthen the information, updating the modal base with the antecedent. In all the best worlds where the miners are in A we block A, and in all the best worlds where they are in B we block B. But since all accessible worlds are either in-A or in-B worlds, it follows that all the best accessible worlds are either block-A or block-B worlds, so none can be block-neither worlds, which falsifies (18d) on the ordering semantics.[40]

The fundamental issue here concerns *nonmonotonicity*, as Nate Charlow (2013) observes.[41] Note first that *simple* necessity or 'all' conditionals are (left downwards) monotonic; i.e. *if p then necessarily q* entails *if r then necessarily q* for any *r* that entails *p*. For example, from (19) we can safely infer (20):

(19) If Garfield is a cat, then he must be a mammal.

(20) If Garfield is a yellow cat, then he must be a mammal.

The ordering semantics does allow for nonmonotonicity, but in only the following way: eliminating the *p* worlds which were initially ranked highest can result in the highest-ranked worlds that remain

---

[39] Due to Regan 1980.

[40] For objections to simple solutions, see Kolodny & MacFarlane 2010.

[41] Koldony & MacFarlane attribute it to the failure of modus ponens, but Charlow proves that the same problem arises for theories of conditionals that invalidate modus ponens.

being *not-p* worlds. However, this explanation is unavailable for the miner puzzle. No worlds are eliminated in the state of ignorance, and apparently all the best worlds are block-A or block-B worlds. What the ordering semantics needs to explain is how adding information can change the *ordering* of possible worlds: what Kolodny & MacFarlane label "serious information dependence". This is a serious challenge, because in the Lewis-Kratzer framework, adding information merely eliminates possible worlds from the domain, which cannot by itself change the relative order of any two worlds within the domain.

The various responses recently offered on behalf of ordering semantics fall into two categories, which both encounter serious difficulties. One (conservative) strategy appeals to an information-sensitive (or subjectivized) ordering source, such as: *in view of what maximizes expected value*.[42] Plausibly speakers can and do sometimes make normative claims about what ought to be done to maximize some salient kind of expected value, and this would indeed generate different orderings as information strengthens. But as a general interpretation of information-sensitive claims, this analysis is problematic. It's natural to suppose that the salient ends or ideals for these claims remain the objective ends (e.g. of saving miner 1, miner 2, etc.), rather than higher-order or subjectivized ends.[43]

This is easier seen in simple cases involving a single end, like Theseus's above. (Or equivalently, suppose blocking neither shaft instead provides a 95% chance of saving all ten miners). Suppose Theseus acquires additional information reducing the likelihood of success if he selects door A from .5 to .25. Given this strengthening of his information, he ought to choose door B (.3 probability of success) rather than door A. But it's natural to say,

> (21) In order to *save his compatriots*, Theseus ought to choose door B.

just as it's also natural to say, relative to his earlier information, that in order to save his compatriots, he ought to choose A. Here we have serious information dependence relative to a constant, *objective* end.

Further evidence comes from data about appropriate modal force. In the (unstrengthened) subjective Theseus scenario, the following is true:

---

[42] von Fintel 2012, Dowell 2013, Silk 2013.
[43] Cf. Charlow 2013, Silk ms.

(22) In order to maximize expected value, Theseus *must* choose door A, and *can't* choose B.

But it's intuitive to say rather that he merely *ought* to choose A, and *could* choose B, as in (17) above, indicating that these information-sensitive 'ought' claims aren't relativized to a subjectivized end like *to maximize expected value*.

Relatedly, Charlow (2013: 2305) observes that serious information dependence seems to arise uniquely for "weak necessity" verbs like 'ought', but not for 'must', discovering intuitions that the following set of sentences aren't consistent:[44]

(23a) If the miners are all in A, then we must block A.

(23b) If the miners are all in B, then we must block B.

(23c) The miners are either all in A, or they are all in B.

(23d) We must block neither shaft.

The conservative strategy isn't subtle enough to accommodate this difference; strengthening information changes what *must* be done in order to maximize expected value. This evidence tells against an information-sensitive ordering source.

The second strategy proposes, more radically, to solve this problem for ordering semantics by appealing to additional (third, fourth…) parameters, and a wider palette of modal operations on these. These solutions are complicated and varied, but encounter problems of their own. Observing the limitation of serious information dependence to "weak necessity", Charlow (2013) follows von Fintel & Iatridou in adopting a secondary ordering source, and proposes that 'ought' semantically allows a "merging" operation between the two ordering sources g_1 and g_2. As he implements this idea to address the miner puzzle, the higher-order ideals in $g_2(w)$ select from the primary ideals in $g_1(w)$ a set of "practical ends" which are the *strongest actionable weakenings* of the primary ideals. An ideal is "actionable" just in case there is an action guaranteeing that outcome relative to the information. In the original miner scenario, the strongest actionable weakening of the ideal *saving all ten miners* is the practical end, *saving at least nine miners*. In view of this practical end, blocking neither shaft is the best option given the limited information.[45] Serious information

---

[44] These intuitions are disputed (e.g. Carr ms.), I think because in the original miner puzzle it's too easy to accommodate context shifts in ends/ideals; see Section 6.

[45] A problem: not *all* worlds where at least nine miners are saved are block-neither worlds. Charlow suggests one solution (2013: 2318-9n): the set of practical ends generated by the merging also includes the negation of any *non-*

dependence is explained by the possibility that strengthening information causes additional or stronger practical ends to become actionable.

Charlow's proposal creates new difficulties. First, it falsely predicts that in the original miner scenario, the following sentences are both consistent (and true) without any shift in context:

(24) We ought to block neither shaft.
(25) #We must block either shaft A or shaft B.[46]

Second, absent a special, as yet untold story about the semantics of the antecedent, it would seem to predict that (26) is true, since saving nine is the strongest actionable weakening of the designated primary ideal of saving all ten:

(26) #(If we want) to save all ten miners, we ought to block neither shaft.

But blocking neither is the only option we may not choose in order to save all ten.

Fabrizio Cariani, Magda Kaufmann and Stefan Kaufmann (2013) propose an additional contrast-set or "decision problem" parameter *d*—as in my own semantics for 'ought'—instead of a secondary ordering source, consisting in mutually exclusive "chooseable" options.[47] This partitions the *f*-accessible worlds into sets defined by chosen option. The ordering source *g* ranks each *partition* by just the ends true throughout it, i.e. those ends guaranteed by the option. Since saving nine miners is the only relevant end guaranteed by any option, worlds are ranked only by their closeness to this end. Serious information dependence is explained by the possibility that strengthening information may lead to additional ends being guaranteed by some option (e.g. given that the miners are in A, blocking A is guaranteed to save ten.) This avoids the problems observed for Charlow, but shares a further problem: appeal to "guaranteed" or "actionable" outcomes is artificially restrictive, because sometimes we ought to pursue a *chance* of an outcome. Even if blocking neither shaft provides only

---

actionable primary ideals; i.e. *not* (*all ten miners are saved*). However, suppose that if we block neither shaft there is a 5% chance the tenth miner survives, while another option guarantees saving only nine. This analysis then predicts we ought not to block neither shaft—because we might thereby save all ten miners.

[46] Charlow suggests this is merely infelicitous, but it seems correct to say, 'It's *not true* that you must block either A or B; you could and should block neither.'

[47] See also Charlow 2013, Carr ms.

a 95% chance of saving nine miners, it plausibly ought to be chosen over a 50% chance of saving everyone (Carr ms.).[48]

A variation from Graham Katz, Paul Portner, and Aynat Rubinstein (2012) offers a fix for this problem. They propose separate epistemic (g_e) and deontic (g_d) ordering sources, and a merging operation on these. As applied to the miner puzzle, the idea is that valued outcomes are first ordered into tiers by how likely (actionable) they are, by g_e, and then are ordered within each tier by their closeness to the deontic ideal g_d. An option *ought* to be chosen iff it corresponds to the deontically best of the epistemically most actionable outcomes. In Jennifer Carr's probabilized miner scenario, this is the option of blocking neither shaft, since the most actionable end (at 95% if one blocks neither shaft) is saving nine. But this fails to accommodate normative judgments that trade low probability for high value. Suppose blocking neither shaft will certainly save exactly one miner, while blocking a shaft gives a 50% chance of saving all ten. Presumably we ought to block a shaft, although saving exactly one is the uniquely most actionable outcome.

I have not shown that serious information dependence is a fatal problem for the ordering semantics. Its champions may reply, as Charlow does (p.c.), that my objections merely target contingent details of the way they have proposed implementing these more general strategies to address particular variations of the miner puzzle. I think this itself reflects a problem: it remains to be shown that any of these strategies can be developed into a concrete theory that systematically generates correct results when applied to arbitrary cases, rather than on an ad hoc basis. In one of the most sophisticated attempts, Cariani (this volume) concludes from these problems that the semantics for 'ought' must be probabilistic, which he implements by modifying his earlier view, above, replacing the modal base with a pair consisting of an information-state and a probability function $<i, Pr>$, and assigning probabilistic content to the premises in the ordering source. I can't do justice to this proposal here, but one problem it faces is explaining why 'must' and 'may' aren't similarly probabilistic, once we've probabilized ordering sources; Cariani suggests a further semantic rule that these terms ignore such premises.

The point I want to make is that here, also, the challenge for ordering semantics is to show that it can generate the same results that my simpler semantics delivers already, without further

---

[48] *Single-Minded Theseus* shows that appeal to a secondary normative ordering source can't be sufficient to explain serious information dependence, since this is a feature also of single-end scenarios; see note 44.

elaboration. For the end-relational semantics provides a straightforward explanation of serious information dependence. Nonmonotonicity is a basic characteristic of probability and 'most' conditionals: strengthening information often switches *most p* from true to false, unlike *all p*. While most philosophers already born are now dead, for example, most philosophers already born *after 1973* are not now dead. Similarly, if the information is strengthened in a way that eliminates the possibilities where Theseus both chooses door A and saves his compatriots, then in most of the remaining possibility-space where his compatriots are saved, he takes door B. (Similarly for the simplified miner case where blocking neither shaft is 95% likely to save all the miners).

Unlike the kind of nonmonotonicity provided by ordering sources in the basic Lewis-Kratzer semantics, the 'most' analysis explains serious information dependence. It also avoids all the problems observed above for versions of ordering semantics. It invokes the same (objective) ends as objective normative claims, accommodating the use of restrictors like *in order to save his compatriots*, and it makes the right predictions about appropriate modal force. It correctly predicts nonmonotonicity uniquely for "weak necessity" verbs like 'ought' and 'should', and not for 'must'. It accommodates probabilities with ease. As already observed, it doesn't imply, counterintuitively, that worlds where agents acts as they subjectively ought are always the "best" worlds : the worlds where Theseus saves his compatriots are best, even if they aren't those in which he acts as he ought. Finally, it predicts the right verdicts in at least these simple single-end cases, because so long as the end is held fixed, changes in expected value just are changes in the conditional probability of the end. The information-sensitivity of 'ought' therefore provides further support for abolishing ordering sources.

However, it may fairly be objected that celebration is premature, as I haven't yet shown that my semantics yields the right results for the original miner puzzle. A full solution to this requires a solution for cases involving multiple ends/ideals, to which I now turn.

6. Conflicts

One argument for an ordering source parameter remains: the need to accommodate inconsistent premise-sets, or conflicts either between incompatible ends/ideals or between ideals and facts. Ordering sources provide a way of reaching a consistent domain from inconsistent inputs. But despite this advertised virtue, here too the ordering semantics encounters a problem: the Lewis-

Kratzer ordering mechanism seems too crude to handle all but the simplest of normative scenarios. An ordering source g ranks a possible world w_1 as better than a world w_2 iff w_1 satisfies all the ideals in the premise set $g(w)$ that w_2 does, plus at least one more. This presents some obvious difficulties when we consider ordinary use of normative modals.

First, it doesn't allow comparisons between two worlds if each satisfies some ideal the other doesn't, although this is an ordinary situation where it is often perfectly clear that one possibility is better than the other. Suppose you have two desired ends, (e_1) attending your daughter's wedding, and (e_2) watching the football match live; here (27) is surely true:

(27) You ought to attend your daughter's wedding rather than watch football.

It is obviously because of the greater *importance* of attending the wedding that it ought to be chosen over watching football. But ordering sources don't encode the degrees of importance (or normative "weight") of different ideals, treating every ideal in the premise-set $g(w)$ on a par. A second problem arises from the interaction of importance with *uncertainty*, as the ordering semantics doesn't allow the intrinsic importance of ends in $g(w)$ to be balanced against the probability or expectation of achieving them, as modeled in decision theory.

Many creative solutions have been proposed on behalf of ordering semantics. We might try capturing differences in importance by assigning ends to higher- or lower-ranked ordering sources. However, this can't capture a properly decision-theoretic framework, in which less important ends can outweigh more important ends either by combining their normative weight or by their greater expected (probability-weighted) value. In Section 5 we observed proposals to "merge" separate ordering sources (either normative and epistemic, or first- and second- order normative), and for probability-sensitive ordering sources. But no general account is offered of how to generate decision-theoretic verdicts systematically, balancing importance against probability in intuitive ways. I think these strategies are unlikely to succeed, for the reason that ordinal scales can never fully reproduce cardinality.

There is a way the ordering semantics can generate any desired ordering of worlds, however.[49] Suppose there are three possible outcomes, to be ranked in the order o_1, o_2, o_3. This result can be

---

[49] E.g. Swanson 2008.

generated by the premise set $\{o_1, (o_1 \vee o_2), (o_1 \vee o_2 \vee o_3)\}$. But this technical fix is problematic as an account of natural language semantics. In the Lewis-Kratzer framework, the ordering $\leq_{g(w)}$ is generated by a semantically fixed rule from the premise-set $g(w)$, which is in turn generated by the conversational background g, a function from the context determined by the speaker's intentions. So the speaker's intended set of ideals is in the driver's seat. By contrast, this solution reverse-engineers a premise-set from the desired ordering, which can't be generated from any intuitive inputs such as degrees of importance and probability. The only kind of conversational background that could plausibly generate such a premise-set is something like: *given the set of ideals that by the semantic rule ≤ yields the ordering corresponding to expected value by the salient probability and value functions*. But it is implausible that this is the kind of conversational background operative in ordinary normative thought. Whereas Lewis-Kratzer ordering sources are usually described in highly intuitive ways (e.g. *in view of what you desire*), on this fix they and their derivative premise-sets bear little resemblance to intuitive inputs for normative thought even in mundane cases.

In view of such problems, many writers reject ordering semantics in favor of *quantitative* ("fully Bayesian") semantics that instead have parameters for at least (i) a *probability* distribution over worlds, and (ii) a *value* function assigning normative weights to outcomes, and give 'ought' the modal force of a decision rule on expected value.[50] This has the virtue of systematically generating the verdicts of decision theory, but there are problems here too. First, whereas an ability to *accommodate* decision-theoretic judgments is a desideratum, such judgments shouldn't be guaranteed true by semantics alone, since competent speakers can vary in their tolerance of risk.[51] Carr (ms.) therefore proposes an additional (third) parameter for a *decision rule*. Second, it isn't clear how to apply these semantics to ordinary epistemic or circumstantial uses of 'ought',[52] which don't seem to involve value functions or decision rules at all.

I believe that the simple dyadic semantics has been abandoned too quickly in the face of this issue, because it supports an attractive solution which makes the extra semantic machinery of ordering and quantitative semantics gratuitous. This solution appeals only to resources that play a central role in Kratzer's own application of the ordering semantics to different uses of modal verbs:

---

[50] E.g. Goble 1996, Wedgwood 2006, Lassiter 2011.

[51] Kratzer 2012: 25, Carr ms.

[52] For the distinction between epistemic and circumstantial *ought*, see Copley 2006.

i.e. to the *pragmatics* (or "pre-semantics") of how the appropriate conversational backgrounds are fixed and premise-sets constructed:

> The way we understand a particular occurrence of a modal can be at least partly explained by an interaction of independently motivated semantic and pragmatic principles. (Kratzer 1981: 61-2)

In particular, it endorses Kratzer's appeal to the *Rule of Accommodation*:

> If the utterance of an expression requires a complement of a certain kind to be correct, and the context just before the utterance does not provide it, then ceteris paribus…a complement of the required kind comes into existence. (1981: 62)

This solution (explored in greater detail in Finlay 2014, Chapter 6) is as follows.

First, given that modal verbs function logically as quantifiers over possibilities, modal claims relative to an inconsistent modal base will generally be pointless. Informally, a modal base provides a partial sketch of some way the world could be, but the world couldn't be such that p and $\sim p$ are both true. On a dyadic semantics, competent interpreters can therefore be expected to accommodate modal utterances using pragmatic cues to identify the speaker's intended, *consistent*, modal base f.

Second, on the end-relational theory, normative modals are relativized to a single end or outcome e in the modal base. When this end isn't articulated explicitly, we can expect it to be the most salient end in the context, which will generally be the relevantly *most preferred/ important* end of the person(s) whose perspective is salient—whether this is the agent, the speaker, the audience, or some further party.

Even putting aside conflicts of ends, this proposal might seem ill-equipped to handle all-things-considered practical contexts involving multiple ends. In the original miner scenario, deliberation can't be concerned with any single end (like saving all the miners, or saving the most we can, etc.), but rather with a set of ends such as {saving miner 1, saving miner 2, …, saving miner 10}.[53] But there is a simple, intuitive solution. In practical contexts (of deliberation or advice), the most salient end can be identified (perhaps only de dicto) as the *overall outcome* that is the relevantly preferred object of pursuit. Call such an outcome the *contextually preferred* end. If the dyadic semantics is

---

[53] In Kratzer's (borrowed) slogan, we need *commas* rather than *ampersands* (2012: 20).

correct, then an ordinary normative utterance of '*s* must/ought to/may do A', simpliciter, can therefore be interpreted by default as asserting that *given f: (it is the case that i, and it is going to be the case that e), in all/most/some of the remaining possibilities, s does A*—where *i* is the salient circumstances or information (with symmetry of choice) and *e* is the contextually preferred end. This provides intuitive resolutions to the puzzles about conflicts.

First, in the case of contrary-to-duty obligations, any end or ideal that is incompatible with the salient circumstances or information *i* will not be (part of) the contextually preferred end *e*, since we don't seriously pursue ends we're sure are unattainable. Therefore, unattainable ends would typically be omitted from the modal base, just as ability modals ignore information about agents' psychological dispositions in Kratzer's analysis (1981: 53). When judging whether you ought to apologize for causing offense, for example, any possibilities where you never violated offense norms are irrelevant. On the other hand, if we tell somebody locked in jail that *if he wants to go to Harlem he could take the A train*, he will immediately assume that his incarceration is being excluded from the modal base; this resolves our case of unattainable goals from Section 2.

Second, whenever a person believes that two ends *e₁* and *e₂* are mutually incompatible, her contextually preferred end won't incorporate both, and for the same reason: we don't pursue unattainable outcomes. Instead, a speaker's contextually preferred end will be roughly the most preferred combination of ends she believes attainable. If you desire to visit your grandmother in Harlem, and your grandmother in Hoboken, but can't do both, then we'd typically advise you best by telling you what you ought to do in order to visit the grandmother you'd most contextually prefer to visit.

Third, in cases where multiple ends are weighed under uncertainty, an agent's overall motivation will balance intrinsic preference for outcomes (importance) against their expectability (probability). The contextually preferred end therefore needn't be intrinsically preferred, or what the agent most desires. In the miner scenario, for example, our intuitive judgments plausibly reflect a contextual preference for *saving at least nine* (whether at 100% or 95% odds) over *saving ten* (at 50% odds).[54] This accounts for the truth of (19d): if we're going to save at least nine miners, it is most likely that we block neither shaft ('We ought to block neither'). In support of this analysis, observe that if asked,

---

[54] In complex scenarios an agent's contextual preference is plausibly a conjunction of probability-weighted outcomes; e.g. *to have a 30% chance of e₁ and a 50% chance of e₂…*; see Evers 2013, Finlay 2014: 162.

'*Why* ought we to block neither shaft?' it's natural to reply, 'In order to save at least nine miners', but not 'In order to save all the miners'.[55]

On this dyadic semantics, the complex interactions between importance and probability are assigned to the psychology of normative thought rather than to semantics. While more work is needed to establish that these pragmatic resources are sufficient to address every issue arising from conflicts,[56] this explanation has noteworthy advantages over both the ordering semantics and its decision-theoretic rivals. It respects linguists' concerns about the artificiality of building quantitative measures into natural language semantics (e.g. Kratzer 2012: 25), yet also easily accommodates decision-theoretic verdicts without ad hoc semantic maneuvers. This is accomplished without either building a decision rule into semantic competence (thereby respecting the semantic desideratum of normative neutrality), or relativizing ordinary normative speech to a decision rule as an additional conversational background.[57] It therefore employs only the same simple semantic resources needed for epistemic and circumstantial uses, not postulating empty ordering sources or other redundant argument-places. Surprisingly, a dyadic semantics might even provide the best natural language framework for handling cases of conflicts.

7. Metasemantic Issues

My objections to ordering semantics may be thought to stem from an idiosyncratically narrow conception of semantics. I've argued that a dyadic semantics is sufficient in part because certain peculiarities of our use of 'ought' can be attributed to "pragmatics". But whereas I've assumed a

---

[55] Superficially this resembles Charlow's appeal to the "practical end" of saving *nine miners and no more*, but the *most* analysis needn't exclude worlds where all ten are saved, as not among the relevantly "best" worlds.

[56] Perhaps most pressing is the difference between 'must' and 'ought' in these cases. Sometimes we "ought" but don't "have to" choose an option because it promotes an end of greater importance, rather than making any single end more likely. Suppose six miners are in A and four are in B; arguably we then ought to block A but don't have to, as we could block B. But to achieve the (contextually preferred) end of saving six, we *must* block A, and *cannot* block B.

Hypothesis: in basic (single-end) cases, using 'ought' will indicate that *p* is merely the most preferred of possibly multiple acceptable options, whereas 'must' indicates that *p* is the only acceptable option (Finlay 2014: 172). Use of these words may therefore carry generalized conversational implicatures to this effect, which would explain the preference to use 'ought' whenever the end is only marginally preferred over others, and reserving 'must' for communicating that other ends are unacceptable. Use of 'may' or 'could' would pragmatically indicate an end that is merely *accepted* rather than preferred. I also argue (2014: 172-3) that this explains the puzzling impropriety of 'You must, but you won't' (e.g. Copley 2006).

[57] In contrast to Carr ms., for example. One objection to this pragmatics-heavy approach is that because it liberally postulates context shifts, it fails to account for many disagreements and valid inferences. I believe any contextualist semantics faces this issue, and that it is amenable to pragmatic resolution (Finlay 2014: Ch. 8).

narrow conception of semantics as investigating the *conventional meanings* of words (etc.), others understand "semantics" more broadly, as investigating the *truth-conditions* of sentences as uttered in contexts, and understand "pragmatics" more narrowly, as concerning what is merely implicated by utterances. Kratzer might seem to endorse this broader conception of semantics by beginning her textbook on compositional semantics by writing that "a theory of meaning…pairs sentences with their truth-conditions" (Heim & Kratzer 1998: 1). Such a semantics for 'ought' would aim not just to identify the "meaning" of the word 'ought' itself, but rather to give a complete theory of the truth-conditions of 'ought' sentences as uttered in contexts.

My claim to have offered a *simpler* semantics may therefore be thought confused, on the grounds that I've really just relabeled what is functionally still an ordering source as part of the "pragmatics". The ordering source's function in the Lewis-Kratzer semantics is nothing other than to generate consistent premise-sets from inconsistent inputs, and it is therefore just a formal implementation of the same kinds of contextual pressures I've described in Section 6. Then I haven't really shown that we can do without ordering sources. Since the ordering source mechanism is flexible enough to yield any ordering we might want, my objections to it as crude also fail.

A theory of the meaning of 'ought' should provide more than this, however. If the task of semantics is just to give a formal statement of the truth-conditions of sentences, then semantic theories would be *nonexclusive*. Kratzer provides one way of modeling the truth-conditions of 'ought' sentences, but these could be modeled with equal legitimacy either by (i) more fine-grained semantic theories with three or more parameters, which separate different kinds of inputs to orderings such as utilities, probabilities, and decision-rules (as in Carr ms.), or by (ii) coarser semantic theories with just one parameter, identifying the background *f* with an often complex function from context which itself eliminated inconsistencies—as in the dyadic semantics I've championed.[58]

This fails to provide other things that semanticists, including Kratzer, clearly want from a semantic theory. First, it has nothing to say about whether a term is lexically ambiguous—if it doesn't eliminate the possibility of lexical ambiguity altogether.[59] Capturing truth-conditions with a single complex formula provides no evidence that a word has a single common meaning—especially once we begin positing parameters commonly left empty. A single rule can be provided even for

---

[58] Hall ms. makes this point forcefully.
[59] Cf. Kratzer 2012: 4.

paradigms of ambiguity like 'bank' and 'mole', by positing a parameter that toggles to different values. There would be no difference between saying that 'must' has a unified meaning involving two parameters including an ordering source which is empty in alethic uses, and saying that 'must' has two meanings, being either a simple dyadic operator, or a polyadic operator taking ordering sources.

Second, it needn't reveal anything about the *semantic competence* of ordinary speakers, and so wouldn't realize Kratzer's stated aim of identifying the "abilities of a person who has a complete grasp of the modal system" of some natural language (1981: 72). Unity is a desideratum of lexical semantics mainly because of the need to explain how speakers are able to learn the language and construct and interpret arbitrary new sentences. A theory of meaning for 'ought' therefore owes us a compositional account of the contribution that competent speakers implicitly understand the word 'ought' itself making to the truth-conditions of a sentence, alongside the contributions of other words and contextual features.

The modal base and ordering source aren't merely inputs to truth-conditions in the Lewis-Kratzer semantics, but *conversational backgrounds*: functions fixed by what the speaker has in mind and made explicit by expressions like 'given what we know' and 'in view of what the law prescribes'. The modal force of 'ought' is the rule that competent speakers apply in identifying the truth conditions of an 'ought' sentence relative to these backgrounds. Different semantic theories are therefore not just formal or notational variants, and it is a serious challenge to the feasibility of the ordering semantics if, as I've argued, it requires "conversational backgrounds" that don't capture what ordinary speakers have in mind, or posits a semantic rule that doesn't reflect ordinary modal or normative thought. If I am right about how ordinary users of normative 'ought' select the relevant domain—by a complex function f picking out both circumstances/ information under symmetry of choice and a contextually preferred end—then the ordering semantics are wrong.

Recent semantic theorizing about modals has fixated on the ideal of unity, but has largely ignored the ideal of *simplicity*, which is also important with respect to the acquisition, use, and interpretation of expressions in natural language. For simple tasks we can expect use of simple tools; to butter a slice of bread we reach for a simple butterknife, not for a Swiss Army knife. We should therefore be skeptical of a theory—no matter how unifying—which claims that when someone says (e.g.) 'If this apple costs less than a quarter, then it must cost less than a dollar', by 'it

must cost less than a dollar' she means (without the usual simplifying gloss making the Limit Assumption) that *for every accessible world u there is an accessible world v that comes at least as close to the ideal determined by the (empty) ordering source g, and for every accessible world z, if z is at least as close to that ideal as v, then this apple costs less than a dollar in z* (Kratzer 2012: 40).

If some modal judgments require more complex semantic resources then we should expect to find special terms to have developed for those purposes, rather than all our modal vocabulary developing extra machinery left unused in simpler applications, like empty ordering sources or modal bases. However, we shouldn't be surprised if instead we find people using simple but flexible tools in creative ways to accomplish more complicated tasks, as I suggested above that pragmatic pressures identify consistent premise-sets in cases of conflicts. (Butterknives can be put to a wide variety of uses!) The dyadic semantics respects this constraint. It doesn't posit any kind of argument for modal verbs which is ever left empty (except in bare logical modalities), and rather than building complex but inflexible rules for reaching consistent premise-sets into the semantics itself, it merely extends Kratzer's appeal to "the interaction of independently motivated semantic and pragmatic principles" (1981: 62).

The fact that normative modalities are lexically the most promiscuous is therefore further evidence in support of the simpler semantics. In addition to dedicated normative verbs, like 'needs to', we can get normative readings of virtually any modal auxiliary, including epistemic verbs like 'might',[60] and ability/ dispositional verbs like 'can' and 'could'. If the ordering semantics were correct, we should rather expect normative modalities to be among the most lexically specialized and discriminating, seeing as they require the greatest number of distinct parameters. This promiscuity is however predicted by my end-relational theory that a normative modality is generated simply by including an end in the modal base.[61]

---

[60] The existence of such readings of 'might' is often denied (e.g. Portner 2009: 50), but consider the normative flavor of suggestive advice in 'To go to Harlem, you might take the A train'.

[61] Our lack of grammatical devices for distinguishing conversational backgrounds as modal bases or ordering sources ('given…', 'in view of…', 'in light of…' etc. are all neutral) is also suggestive, as we should expect linguistic resources to have developed to mark such differences in logical role.

8. Conclusion

The difficulties for the Lewis-Kratzer ordering semantics from the behavior of normative modals, despite stimulating ever more complicated fixes in the literature, can all be avoided in simple and intuitive ways by instead giving up on ordering sources and returning to a simpler dyadic framework. Instrumental conditionals, the "weak necessity" and serious information dependence of 'ought', the information-sensitivity of subjective normative claims, and even judgments weighing conflicting ends under uncertainty all seem more amenable to the simpler dyadic semantics. While more work is needed to show that the dyadic approach is viable, these results provide compelling reasons for believing that the move to ordering semantics was a mistake. Of course, given the desideratum of a fully unifying theory, the data to which semantic theorizing must answer extends far beyond the behavior of English modal auxiliaries in normative uses that we've examined here. But my suspicion is that other kinds of use (e.g. epistemic, counterfactual) raise parallel problems for ordering semantics and will yield to parallel solutions in a dyadic framework.[62]

References

Björnsson, Gunnar, and Stephen Finlay (2010), "Metaethical Contextualism Defended", *Ethics* 121 (1): 7-36.

Boghossian, Paul (2006), "What is Relativism?", in P. Greenough & M. Lynch (eds.), *Truth and Relativism*. Clarendon Press: 13-37.

Cariani, Fabrizio (2013a), "Ought and Resolution Semantics", *Noûs* 47(3): 534-58.

Cariani, Fabrizio (2013b), "Epistemic and Deontic *Should*", *Thought* 2(1): 73-84.

Cariani, Fabrizio (this volume), "Deontic Modals and Probabilities: One Theory to Rule Them All?"

Cariani, Fabrizio, Magda Kaufmann & Stefan Kaufmann (2013), "Deliberative modality under epistemic uncertainty", *Linguistics and Philosophy* 36(3): 225-59.

Carr, Jennifer (ms.), "Subjective Ought".

Charlow, Nate (2013), "What We Know and What To Do", *Synthese* 190(12): 2291-323.

Condoravdi, Cleo (2002), "Temporal Interpretation of Modals", in D. Beaver, S. Kaufmann, B. Clark, L. Casillas-Martinez (eds.), *The Construction of Meaning*. CSLI Publications: 59-87.

Condoravdi, Cleo & Sven Lauer (ms.), "Preference-Conditioned Necessities".

Copley, Bridget (2006), "What Should *Should* Mean?" CNRS/Université Paris 8.

Dowell, Janice (2013), "Flexible Contextualism about Deontic Modals", *Inquiry* 56(2-3): 149-78.

Evers, Daan (2013), "Weight for Stephen Finlay", *Philosophical Studies* 163(3): 737-49.

Finlay, Stephen (2009), "Oughts and Ends", *Philosophical Studies* 143(3): 315-40.

Finlay, Stephen (2010a), "What Ought Probably Means, and Why You Can't Detach It", *Synthese* 177(1): 67-89.

Finlay, Stephen (2010b), "Normativity, Necessity, and Tense", in R. Shafer-Landau (ed.), *Oxford Studies in Metaethics* 5: 57-85.

Finlay, Stephen (2014), *Confusion of Tongues*. Oxford University Press.

von Fintel, Kai, & Sabine Iatridou (2005), "What to Do If You Want to Get to Harlem". http://mit.edu/fintel/www/harlem-rutgers.pdf.

von Fintel, Kai & Sabine Iatridou (2008), "How to Say *Ought* in Foreign", in J. Guéron & J. Lecarme (eds.), *Time and Modality*. Springer, pp. 115-41.

von Fintel, Kai (2012), "The Best We Can (Expect to) Get?", online publication.

Foot, Philippa (1972), "Morality as a System of Hypothetical Imperatives", *Philosophical Review* 81(3): 305-16.

Goble, Lou (1996), "Utilitarian Deontic Logic", *Philosophical Studies* 82(3): 317-57.

Hacquard, Valentine (2011), "Modality", in C. Maienborn, K. von Heusinger, P. Portner (eds.), *Semantics: An International Handbook of Natural Language Meaning*. De Gruyter: 1484-515.

Hall, Keith (ms.), "Normative Neutrality and Kratzer's Semantics for Deontic Modals".

Heim, Irene & Angelika Kratzer (1998), *Semantics in Generative Grammar*. Blackwell.

Huitink, Janneke (2005), "Anankastic Conditionals and Salient Goals", in E. Maier, C. Bary, J. Huitink (eds.), *Proceedings of Sinn und Bedeutung* 9: 140-54.

Jackson, Frank (1985), "On the Semantics and Logic of Obligation", *Mind* 94(734): 177-95.

Katz, Graham, Paul Portner & Aynat Rubinstein (2012), "Ordering Combination for Modal Comparison", *Proceedings of SALT* 22: 488-507.

Kolodny, Niko & John MacFarlane (2010), "Ifs and Oughts", *Journal of Philosophy* 107(3): 115-43.

Kratzer, Angelika (1977), "What 'Must' and 'Can' Must and Can Mean", *Linguistics and Philosophy* 1(3): 337-55.

Kratzer, Angelika (1981) "The Notional Category of Modality", in H.J. Eikmeyer & H. Rieser (eds.), *Words, Worlds, and Contexts.* De Gruyter, pp. 38-74.

Kratzer, Angelika (1991), "Modality", in A. von Stechow & D. Wunderlich (eds.), *Semantics: An international handbook of contemporary research.* De Gruyter, pp. 639-50.

Kratzer, Angelika (2012), *Conditionals and Modals.* Oxford University Press.

LaPlace, Pierre-Simon (1814), *Essai philosophique sur les probabilités.* Courcier.

Lassiter, Daniel (2011), *Measurement and Modality.* PhD Thesis, NYU.

Nissenbaum, Jon (2005), "Kissing Pedro Martinez: (existential) anankastic conditionals and rationale clauses", in E. Georgala & J. Howell (eds.), *Proceedings of Semantics and Linguistic Theory XV.* CLC Publications: 134-51.

Portner, Paul (2009), *Modality.* Oxford University Press.

Regan, Donald (1980), *Utilitarianism and Cooperation.* Oxford University Press.

Rubinstein, Aynat (2013), "Figuring out what we ought to do: the challenge of delineating priorities", *U. Penn Working Papers in Linguistics* 19(1): 169-78.

Sæbø, Kjell Johan (2001), "Necessary Conditions in a Natural Language", in C. Féry & W. Sternefeld (eds.), *Audiatur Vox Sapientiae: A Festschrift for Arnim von Stechow.* Akademie: 427–49 .

Silk, Alex (2013), "Evidence-Sensitivity in Weak Necessity Deontic Modals", *Journal of Philosophical Logic*, online publication.

Silk, Alex (ms.), "'Ought' and 'Must': Some Philosophical Therapy", in Anca Chereches (ed.), *Semantics and Linguistic Theory* 22: 43-64.

Sloman, Aaron (1970), "Ought and Better", *Mind* 79(315): 385-94.

Snedegar, Justin (2012), "Contrastive Semantics for Deontic Modals", in M. Blaauw (ed.), *Contrastivism in Philosophy: New Perspectives.* Routledge: 116-33.

von Stechow, Arnim, Sveta Krasikova, & D. Penka (2006), "Anankastic Conditionals Again", in T. Solstad, A. Gronn, & D. Haug (eds.), *Festschrift for Kjell Johann Sæbø.* Oslo: …

Swanson, Eric (2008), "Modality in Language", *Philosophy Compass* 3(6): 1193-1207.

Swanson, Eric (2011), "On the treatment of incomparability in ordering semantics and premise semantics", *Journal of Philosophical Logic* 40: 693-713.

Wedgwood, Ralph (2006), "The Meaning of 'Ought'", in R. Shafer-Landau (ed.), *Oxford Studies in Metaethics* 1: 127-60.

Wheeler, Samuel C. III (1974), "Inference and the Logical 'Ought'", *Noûs* 8: 233-58.

Wheeler, Samuel C. III (2013), *Neo-Davidsonian Metaphysics*. New York: Routledge.

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

Yalcin, Seth (2010), "Probability Operators", *Philosophy Compass* 5(11): 916-37.