

# ‘Ought’ and Resolution Semantics.

Fabrizio Cariani  
Northwestern University

## 1 Introduction

The standard modal analysis of ‘ought’ has a long-standing flaw. The standard account among semanticists has it that ‘ought’, as it appears in a sentence like,

(1) You ought to clean your room.

should be analyzed as an (appropriately restricted) universal quantifier over worlds. This is part of a larger claim about modality in natural language according to which modality comes in a restricted number of varieties (e.g. possibility and necessity) with context intervening to determine a particular flavor (epistemic, deontic, alethic, etc.). I will refer to the standard view as the *boxing* view, because it treats ‘ought’ as a necessity operator (a ‘box’ in the sense of intensional semanticists). A consequence of the standard view is:

[INHERITANCE] If  $p$  entails  $q$  then  $\Box p$  entails  $\Box q$   
(*notation*: I use ‘ $\Box$ ’ to symbolize ‘ought’.)

One central consideration in favor of INHERITANCE is that it explains a core fact about how ‘ought’ is used in English. You ought to give food to your pets, but not every way of giving food to the pets is something you are permitted to do. You may be under more specific requirements. For example, you ought to give them *non-poisonous* food. Presumably this is not the end of it: you ought to give them non-poisonous food *in decent quantities*, and so on. To put the point in general terms, a semantic theory for ‘ought’ should satisfy the following principle:

[COARSENESS]  $\lceil S$  ought to  $\phi \rceil$  can be true even though there are impermissible ways of  $\phi$ -ing.<sup>1</sup>

INHERITANCE cleanly explains COARSENESS: the connection between a less specific and a more specific *ought*-sentence is entailment. The proposition that you ought to give to your pets non-poisonous food entails the proposition that you ought to give them food.

In contrast, accounts that deny INHERITANCE—we may call them *antiboxing* accounts<sup>2</sup>—are often faulted for their perceived inability to explain COARSENESS. Wedgwood criticizes an anti-boxing proposal by Jackson on these very grounds:

Jackson (1985) has proposed analyzing ‘O(p)’ in counterfactual terms, as meaning, roughly, ‘If it were the case that *p*, things would be better than they would be if it were not the case that *p*’. [...] it seems to me that ‘ought’ is not very well analysed in such counterfactual terms. We often say that something “ought” to be the case when it is very much only a *part* of everything that ought to be the case. (Wedgwood, 2006, p.45, fn 16)

In spite of its apparent plausibility, however, INHERITANCE is subject to important difficulties. Here are three such challenges:

### I. Ross’s Puzzle.

Suppose you and I both accept that Joan paid hefty fees to go to a famous school. Then I may truly utter:

(2) Joan ought to attend her classes.

If INHERITANCE holds, it follows that:

(3) Joan ought to either attend her classes or burn down the philosophy department.

But clearly, this is not always an acceptable inference. (3) communicates the categorical information that Joan has two ways of doing what she ought to; (2) does not. This is *Ross’s Puzzle*.<sup>3</sup> The anti-boxer suggests that the unacceptability of Ross’s inference be taken at face-value. That is, she suggests that, in the relevant context, (3) is false even though (2) is true.

### II. Prof. Procrastinate.

Jackson and Pargetter (1986) discuss this (now famous) case:

PROCRASTINATE. Prof. Procrastinate is invited to review a book on which he is the only fully qualified specialist on the planet. Procrastinate’s notable character flaw, however, is his inability to bring projects to completion. In particular, if Procrastinate accepts to review the book, it is extremely likely that he will not end up writing the review. In the eyes of the editor, and of the whole scientific community, this is the worst

possible outcome. If Procrastinate declines, someone else will write the review—someone less qualified than him, but more reliable.

*J&P*'s intuitions go as follows.

(4) Procrastinate ought to accept and write.

It is false, however, that,

(5) Procrastinate ought to accept.

After all, if he accepts the job, he will not write the review, so the worst possible state of affairs will obtain.<sup>4</sup> If 'ought' satisfies INHERITANCE, this distribution of truth-values cannot be obtained.

### III. Conditional 'oughts'.

On a standard boxing theory (Kratzer 1977, 1981, 1991*b*: details to follow), 'ought' universally quantifies over a set of relevant worlds. The dominant theory of conditionals among semanticists (Kratzer, 1991*a*) implies:

[RESTRICTION] In conditionals of the form  $\lceil \text{If } p, \text{O}q \rceil$ , the conditional antecedent restricts the space of relevant worlds against which the deontic modal in the consequent is evaluated.

Imagine a basketball coach talking to one of her players uttering:

(6) If you have five fouls, you ought to refrain from trying to block shots.

According to proponents of RESTRICTION, the role of the conditional antecedent in (6) is to restrict the worlds against which the 'ought'-sentence in the consequent is evaluated (namely, it restricts to those worlds in which the given player has five fouls).

Together with a boxing semantics for 'ought', RESTRICTION implies the logical validity of the schema  $\lceil \text{If } p, \text{O}p \rceil$ .<sup>5</sup> Among instances of this schema are things like:

(7) If you drink a bucket of poison, you ought to drink a bucket of poison.

(8) If you believe that Rome is in France, you ought to believe that Rome is in France.

Most of us have no doubt that these sentences are false, but the standard semantics (as well as several variations on it) predicts them to be logically true.

Defenders of INHERITANCE resist these challenges by proposing alternative explanations of the data. For Ross's puzzle, they suggest that asserting (3) in a context in which one could also assert (2) would violate Grice's maxim of quantity—and

so we do not need to postulate its falsity to accommodate the salient intuitions. In Procrastinate cases, they argue that the intuitions can be explained away by postulating subtle differences in context. (5) is naturally heard as false because, when we evaluate it, we ignore scenarios in which Procrastinate accepts and writes. In the third class of cases, they respond by either rejecting RESTRICTION or by stipulating some pragmatic constraints that make (7) and (8) inappropriate.

Towards the end of the paper, I criticize these *resistance strategies*, as I will call them. I will argue that they are either insufficient or unsuccessful. One crucial point should, however, be apparent without further elaboration: *collectively* the resistance strategies are unable to capture the unity behind the problematic phenomena. A story that conjoins INHERITANCE with the diverse family of resistance strategies cannot, by design, see any unity in the explanation of the verdicts of unacceptability in (3), (5) and (7).<sup>6</sup>

The dialectical situation, then, appears to be this: INHERITANCE seems at once indispensable to an account of COARSENESS and problematic when taken on its own. A suitable system of resistance strategies is a possible way of resolving this tension. My central task in this paper is to articulate a different, and in my view more compelling, route to the same goal.

I aim to develop an anti-boxing account that takes the intuitions in I-III at face-value as violations of INHERITANCE and still satisfies COARSENESS. This account can be developed so as to depart minimally from currently accepted frameworks for modality in natural language semantics. The benefit of such an account is twofold. On the one hand, we need a better understanding of the anti-boxing position to rationally compare it to the much more established boxing alternative. On the other, I claim to be able to thread through my central desiderata (COARSENESS and a solution to I-III) better than any alternative account. My hope is that, once the theory is laid out, the need for resistance strategies will not feel so urgent.

INHERITANCE-violating semantic theories are no news for deontic logicians. For example, *neighborhood* semantics (see Chellas, 1980, ch. 9, for a canonical application to deontic modalities) was designed to represent the logical properties of certain weak modal logics that do not satisfy analogues of INHERITANCE (called ‘non-normal modal logics’). I take my project to be somewhat different from (though partially overlapping with) that of deontic logic: the job of natural language semantics for deontic modals is to deliver predictions about the significance of speech acts involving *ought*-sentences in a variety of contexts. Intuitions about acceptability of inferences can count as data-points, but the theory must satisfy many further constraints. This contrasts with the theoretical aims of deontic logic, which is just to characterize a class of valid inferences, regardless of how the semantics’ concept of *truth in a model* interfaces with a theory of meaning for natural language. In the sense of Dummett (1973), semantic theories for deontic logics are often *purely algebraic*.<sup>7</sup>

Here is the plan: in §2, I spell out the idea behind the semantics in an informal (framework-independent) way. In §3, I introduce a relatively standard semantic

framework. In §4, I reintroduce my semantics with all the appropriate bells and whistles. In §5, I take up some objections against my anti-boxing approach. In §6, I finally criticize some of the key resistance strategies. This final, negative, section does not presuppose the previous ones: if the reader prefers, she can start with the negative arguments against resistance strategies, and then return to §2.

## 2 The Semantics Informally.

An anti-boxer's first thought is to find a common diagnosis for the violations of INHERITANCE in I-III. Following practice among linguists, I call the proposition embedded in the 'ought' the *prejacent*.<sup>8</sup> In each case, the problematic sentence (i.e. (3), (5), (7)) expresses an 'ought' that seems undermined because some course of action compatible with the prejacent is impermissible relative to the salient norms. For example, in (3), Joan's burning down the philosophy department is (usually) impermissible, and that seems related to the unacceptability of (3). In fact, in bizarre conditionals such as (7), *all* of the possibilities that satisfy the prejacent are intuitively impermissible (in most contexts). One may be inclined to base an anti-boxing proposal on a constraint like:

[FALSEMAKING]  $\lceil S \text{ ought to } \phi \rceil$  is false (in a context) if there is an impermissible (relative to the salient norms, desires, values, etc.) way of  $\phi$ -ing.

It is promising that FALSEMAKING aims to capture a common feature of the examples (furthermore, one which appears to explain our intuition of rejection in some of them). However, FALSEMAKING directly contradicts COARSENESS. The apparatus I present informally in this section is designed to do better.

### Resources.

On any theory, the semantic value of an *ought*-sentence is determined compositionally from the semantic value of the prejacent together with some additional parameters. For example, an unadorned version of the standard boxing theory looks like this:

$\lceil S \text{ ought to } \phi \rceil$  is true in the actual world iff  $\lceil S \text{ does } \phi \rceil$  is true at all of the *deontically ideal* worlds.

According to this proposal, the truth of *ought*-sentences depends on the compositional semantic value of the prejacent and on the value of an external parameter—some ranking of the relevant possibilities.<sup>9</sup>

In this respect, my theory is no different. The semantic value of an *ought*-sentence depends in part on semantic features of its prejacent and in part on external parameters. I do, however, require a more elaborate set of parameters, namely:

- (i) a range  $o_1, \dots, o_n$  of (mutually incompatible) alternative courses of action available to the salient agent. If an action is represented in the relevant range, I call it an *option*.
- (ii) some way of ranking options (I call this the *ordering*).
- (iii) a threshold in the ranking that distinguishes permissible options from impermissible ones (I call this the *benchmark*).

My account does not require any particular view of how these parameters receive their values. It will work regardless of whether the ordering is set by context of utterance, by objective moral norms, by your boss's desires, or by a context of assessment. For determinacy, I assume that all parameters get their initial value set by the context of utterance. I do not give an argument for this and my central conclusions do not at all depend on this.

The contextualist stance, although not mandated, does appear particularly attractive for the class of alternative actions. Indeed, authors who invoke contrast classes typically assume them to be determined by context.<sup>10</sup> I propose a way of developing the contrastive insight explicitly designed to address the problem of COARSENESS. Suppose you and I are discussing our mutual friend Jenny's plans for how to get to school. Consider representing her deliberation as a choice among the alternatives that are in the set:

(\*) {running, walking, swimming, driving}.

Assume also that, given background knowledge, these options are incompatible. (\*) represents a certain level of detail at which we could describe Jenny's deliberation. This does not mean that the options in (\*) are the most specific courses of actions Jenny can adopt. After all she can choose to walk in a blue dress, or walk in a red dress. It does mean, however, that distinctions between various finer ways of performing the actions listed in (\*) are contextually not salient. If we needed to represent such distinctions, we could use a finer set of options:

(\*\*) {running, walking in a blue dress, walking in a red dress, swimming, driving}.

One advantage of this way of thinking is that no metaphysics of basic and complex actions is required: the relevant options need not be metaphysically distinguished by being especially simple types of action.

### The Blueprint.

It is natural, at this point, to ask: given a context and an agent, how do we determine which options are relevant for that agent?<sup>11</sup> In one respect, I would prefer avoiding commitment to such a criterion. As far as compositional semantics goes, this is no more objectionable than refraining to specify a rule that determines the precise extension of 'here' as uttered by a particular speaker in a given context. Nonetheless, it will help, for determinacy's sake, to have at least one natural way of thinking about

the problem. Suppose that context records, for each salient agent, a set of goals, values and desires (as we will see, the standard semantic framework independently assumes something like this). At a first stab, the contextually relevant options for agent  $\alpha$  are the coarsest courses of action that are differentiated by  $\alpha$ 's salient goals, values and desires. If Jenny's goals are the only thing that matters, and Jenny's goal is simply to get to school without sweating and without using gasoline, then (\*) is a more appropriate set of options than (\*\*). If, in addition, she'll get fined for wearing a blue outfit and she wants to avoid this, (\*\*\*) is a better choice of options.

So, let us suppose we do have options and that they are the objects of the ranking in our semantic machinery. A range of options like (\*) (plus the benchmark) may be ranked as follows:

running > walking = swimming > benchmark > driving

with driving being the only impermissible option. For the purposes of this informal sketch, I am going to say that certain actions are *ways of* doing other actions. In the intended sense, driving to school is a way of going to school. I emphasize that this talk is only for illustration purposes and will be avoided in the formal theory.

Not every *ought*-sentence can be evaluated against any set of parameters. For example, the options in (\*) seem unsuited to evaluate:

(9) You ought to have a cup of coffee.

At the same time, those options should allow us to evaluate:

(10) You ought to run or walk.

To this effect, given a set of options, I distinguish two different kinds of prejacents. When an action  $a$  perfectly divides the options in two classes (those options that are ways of performing  $a$ , and those options that are not ways of performing  $a$ ), I say that  $a$  is *visible*. Running or walking is visible in (\*), having a cup of coffee is not. A visible action  $a$  can be:

- *permissible*: some option that is a way of performing  $a$  meets or exceeds the benchmark.
- *strongly permissible*: every option that is a way of performing  $a$  meets or exceeds the benchmark.
- *optimal*: all of the *deontically ideal* options (according to the ranking) are ways of performing  $a$ .

Given the sample ranking above, the action of running or driving is permissible and optimal but not strongly permissible; running is all three, while swimming or driving is permissible, but neither strongly permissible nor optimal.

The informal core of my account of 'ought' is as follows:

$\lceil S \text{ ought to } \phi \rceil$  is true (relative to a set of parameters) iff  $S$ 's  $\phi$ -ing is *visible*, *optimal* and *strongly permissible* (relative to those parameters).

After turning this blueprint into a fleshed out account (§4), I will discuss some possible variations in response to some objections. For now, let me turn to showing how to use the theory to address our initial puzzles.

### Explaining the Puzzles.

It is immediate that this account implies the expected verdicts in the motivating cases. A sentence of the form  $\lceil \bigcirc(p \text{ or } q) \rceil$  may be false even if  $\lceil \bigcirc p \rceil$  is true. This can happen for example if the option specified by  $q$  is visible and impermissible. Consider again:

(3) Joan ought to either attend her classes or burn down the philosophy department.

Joan's options may look like this:

{attend class, stay home, burn down the philosophy department}

I grant that in most contexts, burning down the philosophy department will not be an antecedently relevant option. However, it is certainly natural to assume that an utterance of (3) will *make* that option relevant<sup>12</sup>. It is also natural to rank it below the benchmark (though there might be unusual contexts in which this is not so). If so, (3) can be false while (2) is true.

In the Procrastinate case,

(5) Procrastinate ought to accept.

can be determined to be false provided that the option of accepting and not writing is salient and that it is ranked below the permissibility threshold. That is to say, Procrastinate's options must be no coarser than:

{accept and write, accept without writing, do not accept}

And those options must be ranked as:

accept and write, > do not accept > benchmark > accept without writing

Since, this matches the description of the case I conclude that my account can easily account for the distribution of truth-values.

Finally, let's consider again:

(7) If you drink a bucket of poison, you ought to drink a bucket of poison.

What suffices to explain the falsity of (7) on my view is that 'drinking a bucket of poison' is an impermissible relevant option—it is not only antecedently impermissible, but it is also *conditionally* impermissible. That is, once we update the modal base with the information in the antecedent, the consequent remains impermissible.<sup>13</sup>



## Satisfying COARSENESS.

In virtue of the *optimality* requirement, my proposal shares some features with a boxing semantics. Differences with the boxing account arise from the interaction of two features: the first is that the *relata* of the ordering are not worlds, but rather *options*. The second is the introduction of a permissibility threshold. The combination of these two features allows it to vindicate the anti-boxer's insight without violating COARSENESS: ' $S$  ought to  $\phi$ ' can be true even though there are impermissible ways of  $\phi$ -ing, provided those are not relevant. To make this point vivid, notice that the account satisfies a 'safe' analogue of FALSEMAKING:

[COARSE FALSEMAKING] An *ought*-sentence is false (in a context) if there is a relevant option compatible with the prejacent that's impermissible relative to the salient norms.

This handles the central case I used in motivating COARSENESS. The option of feeding the pets poisoned food is generally not visible: whether it is visible or not is determined by principles governing the dynamics of context.

It is worth comparing my approach with prominent contrastive anti-boxing ideas (e.g. Jackson (1985) and Jackson and Pargetter (1986)). On these views, ' $\bigcirc p$ ' expresses a kind of comparison between the proposition expressed by  $p$  (in context) and one or more contextual alternatives. My approach differs from this for multiple reasons. Conceptually, on my account, the options play a role that is very similar to the role that possible worlds play on the standard account. They are points at which propositions can be deemed to hold or fail to hold (or neither, in the case of non-visible propositions).

As a consequence, given an *ought*-sentence, there may be multiple relevant ways of actualizing its prejacent. On typical contrastive accounts, this is not true. There is only one option corresponding to the prejacent and the remaining options are *alternatives* to it. This feature allows me to address the COARSENESS intuition in the particular way that I just sketched, but also produces one more general consequence. My view requires less fluctuation in the class of alternatives—even in those cases that an anti-boxer should supposedly get right. For example, my semantics always allows evaluating (2) and (3) relative to the same set of alternatives.

(2) Joan ought to attend her classes.

(3) Joan ought to either attend her classes or burn down the philosophy department.

The standard contrastive account does not. Suppose (2) is evaluated against the alternatives: go to class, burn down the department, stay home. Then (3) cannot be evaluated against the same alternatives, but rather something like {go to class or burn down the department, stay home}, in which the first is understood as a disjunctive act

of sorts.<sup>14</sup> This more moderate fluctuation also differentiates my account from contrastive accounts that satisfy INHERITANCE, like Sloman (1970) and more recently Finlay (2009).

To conclude this presentation, it is worth highlighting a promissory note. I intend my apparatus to handle non-deliberative uses of ‘ought’, such as:

(11) The car ought to start up in less than 5 minutes.

The talk of ‘actions’, ‘options’, ‘deliberation’ etc. is inappropriate here, but the rhetoric is not an essential ingredient of the formal account I advance in section 4. I defer to later work to investigate the salient differences.

### 3 Standard Semantic Theories

One of my goals is to develop an anti-boxing alternative that is otherwise as conservative and precise as possible. To give us a benchmark for ‘conservativeness’, I now introduce a plain, boxing semantic theory. Although this section covers boxing views (which I reject), it fixes the general framework I operate in. The next section adapts the framework to an anti-boxing proposal.

The general lines of the framework are relatively conventional *possible-world semantics*. Start with a set of worlds representing logical space. Subsets of this space are *propositions*. To each sentence of the language assign a truth-value, either True or False, relative to a context, a world of evaluation and a variable assignment. We call any triple of context, world, and assignment a *point of evaluation*. Truth at a point of evaluation is defined recursively, on Tarskian lines. Since I mostly avoid discussing quantified sentences, I also avoid cluttering the notation with explicit mention of variable assignments. I use the widespread notation  $\llbracket e \rrbracket^{c,w}$  to denote the semantic value (extension) of  $e$  at the point given by context  $c$  and world  $w$ . Relative to a context, we can associate with each sentence an *intension*: the intension of a sentence is a function from worlds to extensions. I refer to the intension of  $e$  at  $c$  as  $[e]^c$  (the formal definition is  $[e]^c = \lambda w \llbracket e \rrbracket^{c,w}$ ).<sup>15</sup>

Each context  $c$  has associated with it a world  $w_c$ —the *world of the context*. Given this, we may define truth at a context simpliciter:

$p$  is True in context  $c$  iff  $\llbracket p \rrbracket^{c,w_c} = \text{True}$ .

Two different notions of entailment can then be defined: preservation of truth at an arbitrary context-world pair, and preservation of truth at a context.

$p$  *simply entails*  $q$  iff for all  $c, w$ , if  $\llbracket p \rrbracket^{c,w} = \text{True}$  then  $\llbracket q \rrbracket^{c,w} = \text{True}$ .

$p$  *properly entails*  $q$  iff for all  $c$ , if  $\llbracket p \rrbracket^{c,w_c} = \text{True}$  then  $\llbracket q \rrbracket^{c,w_c} = \text{True}$ .

These notions are best kept distinct (for reasons specified in Kaplan 1989), but the difference will not matter to my argument. These definitions can be generalized in the obvious way to entailments between sets of premises and a conclusion.

On top of this structure, we can start thinking about modals. It is a crucial motivational point of Kratzer’s framework, which I am about to introduce, that we should not try to capture the variety of modality in English by postulating different lexical entries for the various interpretations of modality. Rather, the different readings of modals all result from a relatively limited core of lexical entries with contextual parameters set in appropriate ways. How does context contribute to the interpretation of a modal? The beginning of an answer to that question lies in the observation that modals in general seem to have relativizations to *backgrounds* of sorts:

(12) [In view of our evidence] Joseph must be at the movies.

(13) [In view of his desires] Joseph must go to the movies.

These relativizations generally remain implicit, but can, if necessary, be made explicit. Kratzer’s framework accounts compositionally for these implicit relativizations by building these backgrounds as parameters in the semantics—more specifically as restricting the domains of the metalinguistic quantifiers employed in their lexical entries.<sup>16</sup>

Kratzer (1977, 1981, 1991*b*) proposed that context determines *conversational backgrounds* (modeled as *sets* of propositions). One of the functions of conversational backgrounds is to determine a set  $\mathcal{M}_c$  of salient worlds. Kratzer calls the conversational background that plays this job *modal base*. Epistemic modals go with an epistemic modal base (that determines a set of possible worlds that are compatible with some body of information), while metaphysical modals go with a *metaphysical modal base*.<sup>17</sup>

Deontic modals are relativized to at least two backgrounds. One is a modal base: Kratzer’s formalism leaves it open what kind of modal base goes with deontic modals. Within the tradition that has grown out of her work, it is common to assign to deontic modals a *circumstantial* modal base (that is to say, the modal base is composed of a set of *facts*, rather than an epistemic state). However, following Kolodny and MacFarlane (2010), I treat modal bases as *information states* (either of the participants to the conversation or of the agents involved in the action).

The second background keeps track of the relative goodness of the worlds according to whatever norms, desires and values may be contextually salient. In Kratzer’s original semantics, the job of fixing the ordering  $<$  belongs to a conversational background she calls the “ordering source”. Think of the ordering source as some set of propositions (e.g. {Jenny gets to school, Jenny does not sweat, Jenny does not consume gasoline}) that determines a partial ordering  $<$  (the pattern of determination is unimportant for my current purposes). We can use the ordering  $<$  to define a *selection function*  $\beta_c$ , taking a set of worlds  $S$  as input and returning the ‘best’ worlds in  $S$ , according to the ordering.<sup>18</sup>

The semi-official lexical entry for ‘ $\bigcirc$ ’ then is:

(B2)  $[\bigcirc p]^{c,w} = \text{True}$  iff  $\forall w' \in \beta_c(\mathcal{M}_c)[p]^{c,w'} = \text{True}$ .

This is not fully correct, because in Kratzer’s semantics some operators (like conditionals) can shift the modal base against which the modal is evaluated. But it is good enough for our purposes.<sup>19</sup>

As I emphasized, Kratzer’s framework is designed for a boxing account of ‘ought’ (and other modals). I now set out to show that, with a slight enrichment, this framework can become hospitable for anti-boxing proposals built on the blueprint I sketched in the previous section.

## 4 Resolution Semantics

### The Formal Account.

In this section, I show how to implement the ideas of §2 within a variation of the formal framework I just laid out. In §2, I talked about ‘actions’, ‘options’ and such. How can we model these concepts within an intensional semantics?

A common proposal is to model action-types as unstructured propositions of some sort.<sup>20</sup> This is not a point of metaphysical doctrine; rather, it is a modeling choice, which I explain as follows: action-*tokens* are *events*, i.e. concrete particulars. However, besides individual action-tokens, say ‘John’s baking of the bread’, one can introduce various other entities, which could be used to model different notions of ‘action type’. The action-type of ‘baking the bread’ we could model as a function from *worlds, times and agents* to truth-values, i.e. as a property that an agent may have (in a world at a time). If we fix the agent, we obtain a different sense of ‘action type’, e.g. *John’s baking of the bread*: this sense can be modeled as a function from *worlds and times* to truth-values. If we ignore the complication introduced by times, possible-world propositions become useful surrogates of action-types in this second sense.<sup>21</sup>

If individual options are modeled as propositions, a range of mutually exclusive options can be thought of as a set of mutually exclusive propositions—i.e. as a partition of a subset  $S$  of logical space. Outside of  $S$ , there will be worlds where the agent does not act at all (e.g. falls asleep, or gets kidnapped by aliens) as well as worlds in which the agent acts in ways that did not appear in the initial range (perhaps because they were irrelevant). Now, lump all these worlds together in a single-catch-all proposition and add it to the partition of  $S$  and you have a partition of the set of possible worlds—call it  $Q_c$ .  $Q_c$  represents the relevant choices available to the agent (plus the catch-all proposition) in context  $c$ . In a given context, the partition, induces a partition  $Q_c \cap M_c$  of the modal base (the partition whose cells are the intersection of cells of  $Q_c$  with  $M_c$ ). To use a term coined by Yalcin (2009), the sets of alternatives provide a *modal resolution*. In the case of ‘ought’, a modal resolution is a representation of the level of granularity of the salient agent’s choice. Inspired by Yalcin’s terminology, I call the approach I develop here *Resolution Semantics*.

As I anticipated, the key differences with Kratzer’s account are:

- (a) that the ordering  $<$  ranks not worlds but options (hence, mutually exclusive and exhaustive *sets* of worlds.<sup>22</sup>)
- (b) that there is a designated benchmark below which options are impermissible.
- (c) following Kolodny and MacFarlane (2010), the ordering can vary as the information state varies (so we should really write  $<_{\langle c, \mathcal{M}_c \rangle}$ : this is evidently clumsy so I will just write  $<$  leaving the other relativizations implicit).

Just as in the standard semantics, we can use the ordering  $<$  to define a selection function  $\beta$  that selects the best options in  $Q_c$ , relative to information state  $\mathcal{M}_c$  (so ' $\beta_c(\mathcal{M}_c, Q_c)$ ' denotes a set of options).

Let  $[p]^c$  be a proposition and  $o$  be an option. I define:

$$[p]^c \text{ holds at } o \text{ iff } o \subseteq [p]^c.$$

Informally, a proposition holds at an option just in case the option entails the proposition. This is a technical definition that supersedes the talk in §3 of options being 'ways of' performing actions. The example I gave back then was that walking to school is a way of going to school. Indeed, the proposition that the agent is walking to school entails the proposition that the agent is going to school. The definition also helps streamline my lexical entry, which is:

- $\llbracket \bigcirc p \rrbracket^{c,w} = 1$  iff
- (i) for every  $o$  in  $\beta_c(Q_c, \mathcal{M}_c)$ ,  $[p]^c$  holds at  $o$ .
  - (ii) for every  $o$  such that  $[p]^c$  holds at  $o$ ,  $o \geq \text{benchmark}$ .
  - (iii)  $[p]^c$  is visible in  $Q_c \cap \mathcal{M}_c$ .

(i), (ii) and (iii) correspond respectively to Optimality, Strong Permissibility<sup>23</sup> and Visibility in the informal account of §2.<sup>24</sup>

### Significance.

In the remainder of this section I want to make two points concerning the position of my account in the larger theoretical context. First, many interesting views on 'ought' can be pinned down as special cases of my proposal. Perhaps most strikingly, a standard *boxing* semantics turns out to be a special case of my account. This happens if:

- (a) The cells in  $Q_c$  are all singleton sets.
- (b) For every  $o$ ,  $o \geq \text{benchmark}$ .

Clause (a) states that the modal resolution is the finest possible (each cell contains a single possible world). Clause (b) neutralizes the role of the benchmark. Of course, given how I described the apparatus, these conditions seldom make sense for deliberative 'oughts': if we think for example that options must be under the relevant agent's control, then (a) is utterly implausible. No agent has it under their control

to bring about a particular world. The conditions may, however, well arise in the case of non-deliberative ‘oughts’. In particular, if there is an *epistemic* use of ‘ought’, it is compatible with my proposal that it be treated as a box. This also suggests that my account allows recognizing a very strong unity (perhaps even perfect identity, though I doubt it) among the lexical entries of modals with different flavors. To put the point more simply, it is compatible with my account that we do not need to postulate lexical ambiguities to handle the diversity of meanings of ‘ought’.

Second, my account is evidently inspired by Lewis’s relevant alternatives theory of knowledge in his (1996). Lewis tries to reconcile fallibilism with the idea that knowledge that  $p$  requires ruling out all of the alternatives to  $p$ .

Our definition of knowledge requires a *sotto voce* proviso.  $S$  knows that  $P$  iff  $S$ ’s evidence eliminates every possibility in which not- $P$  - Psst! - except for those possibilities that we are properly ignoring. (Lewis, 1996, p.554)

My requirement of COARSENESS is structurally similar to the fallibilist constraint on knowledge and my constraint of COARSE FALSEMAKING is structurally similar to Lewis’s maneuver.

[COARSE FALSEMAKING] An *ought*-sentence is false (in a context) if there is a relevant option compatible with the prejacent that’s impermissible relative to the salient norms.

Having said that, my proposal does not stand or fall with Lewis’s. First, as I mentioned, the idea that deontic modalities require some contrast class is much more entrenched than its analogue for knowledge. Second, unlike Lewis, I do not aim to defend a particular set of rules that determine which options are relevant. Third, as I mentioned, my account does not require being a contextualist about the range of options. For each agent (in a world and at a particular point in time) there is a privileged partition, namely, the partition whose cells are the finest courses of action available to that agent (in that world and at that time).

## 5 Objections.

In this section, I take up two objections that naturally arise against accounts like the one I just developed. Both objections have to do with the logic that results from my account. The first is based on the observation that ‘ought’ and permissibility do not come out as duals. The second is that my account seems to settle some substantive issues on which a formal semantics should just remain neutral. I discuss each in turn.

### Duality.

I started out by observing that INHERITANCE is among the tenets of the Boxing view.

(IN) If  $p$  entails  $q$  then  $\bigcirc p$  entails  $\bigcirc q$ .

This principle is clearly one that an anti-boxer should reject, and it does fail on my proposal. However, it is possible to give a very simple proof of INHERITANCE from premises that appear plausible and that have typically been accepted by writers on deontic modality. Let's use  $[\mathcal{P}]$  to symbolize permission. First, it's standard to think that 'ought' and permission are duals.<sup>25</sup>

[DUALITY]  $\lceil \bigcirc p \rceil$  is equivalent to  $\sim[\mathcal{P}] \sim p$

Second, even if you do not accept the semantics I gave in the previous section, you might accept an analogue of Inheritance for permission.

[PI] If  $p$  entails  $q$ ,  $[\mathcal{P}]p$  entails  $[\mathcal{P}]q$

As it turns out, these two are sufficient to prove INHERITANCE.<sup>26</sup>

Although I reject INHERITANCE, I do not necessarily think we should reject (PI).<sup>27</sup> Fortunately, I think the anti-boxer has a much easier time rejecting the right-to-left direction of DUALITY. Importantly, all of the cases that might persuade one to be an anti-boxer are *ipso-facto* arguments against this principle. Consider again Ross's Puzzle. The anti-boxer insists that (3) is false: it is false that Joan ought to either attend her classes or burn down the Philosophy department. It does not follow from this that she's permitted to do something incompatible with the prejacent of (3) (e.g. go to a museum). She must, after all, attend her classes. The rejection of the right-to-left direction of DUALITY is an immediate consequence of the rejection of INHERITANCE. I am not committed to any failures of DUALITY other than those I have independent reason to accept.

### 'Inconsistent prejacent's'

I am neutral on whether there can be true *ought*-sentences with incompatible prejacent's, but the semantics I described is not. This would be a problem if there was not a ready-made remedy. Such a remedy, however, exists. There is a well known device to turn any semantics similar to mine into one that allows true 'oughts' with incompatible prejacent's.

The device is introduced, among other places, in van Fraassen (1973). Instead of supposing that we have only *one* ordering  $<$ , suppose that we have a set  $\mathcal{P}$  of them. A metaphor to make the device vivid: imagine having a stack of different laws, each of which makes logically consistent prescriptions, but such that, collectively, the laws make logically inconsistent prescriptions. Say then that you ought to do  $\phi$  just in case some law in the stack mandates  $\phi$ .

The point can be fleshed out with precision (and without any reference to laws). Consider the lexical entry I have given for ' $\bigcirc$ '. It shows how to specify the compositional semantic value of  $\lceil \bigcirc p \rceil$  relative to a modal base  $\mathcal{M}$ , a partition  $Q$  and an ordering  $<$ . Call this the *basic definition*. Finally, say that:

$\lceil \bigcirc p \rceil$  is true relative to  $\mathcal{M}, \mathcal{Q}$  and a stack of orderings  $\mathcal{P}$  just in case there is an ordering  $<$  in  $\mathcal{P}$  such that  $\lceil \bigcirc p \rceil$  is true *according to the basic definition* at  $\mathcal{M}, \mathcal{Q}, <$ .

Relative to each particular ordering, true *ought*-sentences must have consistent pre-jacents. But relative to the entire stack of ordering it is possible to have norms with inconsistent pre-jacents.

In conclusion, although the semantics appears not to be neutral on this issue, it can easily be amended according to whatever substantive view of the matter one finds most plausible.

## 6 Resistance Strategies.

I already advertised one advantage of my account over the boxing account: it captures the unity of the recalcitrant phenomena of section 1. The boxer must resort to resistance strategies, and in fact to completely different resistance strategies for each of the problematic data.

In this section, I want to take this debate a step further: I claim that the resistance strategies appear to be incomplete or unsuccessful when taken at face value. I take them up in order, starting with a discussion of the Gricean approach to Ross's paradox. I devote the most space to this resistance strategy because it is the best articulated of the three, and also the one with the fewest alternatives. Conversely, I will have the least to say about the Conditional 'Oughts' problem: the space of possible alternatives is simply too large to be attacked in this context, so I limit myself to a quick remark against one existing proposal.

### The Pragmatic Account of Ross's Puzzle:

Boxers treat the inference from (2) to (3) (both repeated below) as logically valid, but pragmatically unacceptable.

(2) Joan ought to attend her classes.

(3) Joan ought to either attend her classes or burn down the philosophy department.

The intended analogy is with:

(14) Everyone is Italian.

(15) Everyone is Italian or French.

Uttering (15) communicates that the speaker is not in a position to assert (14) and hence it is an odd thing to say if one is in a position to assert (14). Analogously, (3) is an odd thing to say if one is in a position to assert (2). This much is predicted by Grice's maxim of quantity (*Be Informative!*) together with a boxing semantics. Here is how Wedgwood (2006) summarizes the point:



There is an obvious Gricean explanation for why (3) seems an odd thing to say. It is much less informative than something else one might say—namely (2). Asserting the weaker claim would tend to be a useful contribution to a conversation only if one was not in a position to assert the stronger claim.<sup>28</sup>

Additionally, if one can also assume that the speaker is knowledgeable about what Joan ought to do, the implicated content can be taken to be a bit stronger:

IMPLICATED: Uttering (3) communicates that Joan has two ways of doing what she ought to.

If IMPLICATED is indeed ordinarily implicated by an utterance of (3) it would naturally explain the Ross intuitions—why (3) might appear unacceptable in a context in which (2) is acceptable. After all, by asserting (3) one implicates that Joan do as she ought to by either attending her classes or by burning down the department.

The full pragmatic account of how IMPLICATED is triggered would require that the boxer provide much more theoretical detail: it would be a difficult derivation and I doubt that she can carry it out successfully, but my objection is independent of this. My objection is that the way in which IMPLICATED shows up in embeddings is not predicted by the Implicature theory.

Schematically, if an utterance of  $p$  carries a generalized quantity implicature to  $q$ , it is not generally the case that  $p$  can be rejected because  $q$  is false. For example, (15) implicates that not everyone is Italian. As predicted, one cannot reject (15) if it turns out that everyone is Italian. A more natural version of the same test is as follows. Suppose that utterances of  $p$  carry a generalized quantity implicature to  $q$ . You cannot coherently utter: ‘I doubt that  $p$  and  $\sim q$ ’. This test predicts, correctly, that you cannot utter things like:

(16) # I doubt that everyone is Italian or French. In fact, everyone is Italian.

The same pattern holds for *epistemic* modals taking scope over disjunctions. Consider:

(17) # I doubt that Lynn must have either worn a tie or a scarf. In fact, she must have worn a scarf.

This is part of the evidence that the inference from (18) to (19) is a quantity implicature.

(18) Lynn must have either worn a tie or a scarf.

(19) Lynn might not have worn a scarf.

So far so good. The problem is that ‘ought’ does not pass this test. (20) is perfectly intelligible:

(20) I doubt that Lynn ought to either wear a tie or a scarf. In fact, she ought to wear a scarf.

There are two problems here: one for the implicature theory and the other for boxers who want a strong unification of epistemic and deontic modality. The problem for the implicature theory is that *IMPLICATED* fails to be detected in (20) by a principled test that detects quantity implicatures. The problem for the boxer is that there are perfectly parallel cases with epistemic ‘must’ that pass the test.

This behavior is systematic and extends to a range of related environments. Consider Retractions. Suppose that I utter  $p$ , and  $p$  implicates  $q$  via the maxim of quantity (not defeated in the current context). Suppose that I come to learn that  $p$  is true but  $q$  is false. Although we have devices to signal this defect of my original assertion, *outright retraction* is not one of these. Consider:

### **The Bridge Game I**

*Sam*: Joe must have followed suit or played the king of trumps.

*Joan*: Jane had the king of trumps. He must have followed suit.

*Sam*: # I guess what I said was wrong.

To fix terminology, let’s refer to Sam’s first assertion in each of the above cases as the *main claim*, to his second assertion as the *negative reassessment*. My claim is that it is generally inappropriate to respond with a negative reassessment upon learning something that contradicts an assertion’s quantity implicatures but not its content. Disjunctive ‘oughts’, however, can be retracted in this way.

### **The Bridge Game II**

*Sam*: (According to the rules), you ought to follow suit or play the king of trumps.

*Joan*: No, the rules quite explicitly say you ought to follow suit, no matter what.

*Sam*: I guess what I said was wrong.

Sam’s retraction must latch on the falsity of *IMPLICATED*, but it can only do so if it is not an implicature of quantity.

## **Non-boolean Accounts of Disjunction.**

The same sorts of cases, and the same contrast between epistemic modals and deontic modals tells against one last Resistance Strategy. I call this *Blame Disjunction* (BD, for short). BD-supporters believe, as I do, that Ross’s puzzle is due to a genuinely semantic phenomenon. However, they impute it to some semantic fact about disjunction. In particular, they emphasize the similarity between Ross’s puzzle and *Free Choice* effects. It is well known that:

(21) Janet may have cake or ice-cream.

is naturally interpreted as having the truth-conditions ‘Janet may have cake and Janet may have ice-cream’. They hold that a unifying solution to all of these problems consists in a non-boolean account of disjunction.<sup>29</sup>

Of course, any BD-proposal needs to be evaluated in detail, but I limit myself to one observation here.

BD accounts do not predict that deontic modals and epistemic modals should give rise to disanalogous predictions. In fact they naturally predict the opposite—that an epistemic ‘must’ taking scope over a disjunction should pattern in the relevant respects with a deontic ‘ought’ in the same position.

If the arguments I just gave against the pragmatic approach work, they apply with equal force to BD approaches: the data reveal a difference between epistemic and deontic modalities, but the accounts do not reflect this. Hence, the force of Ross’s puzzle is not correctly captured by a BD account.<sup>30</sup> I do not, in principle, rule out the possibility of a BD-account that implies that disjunctions scoping under epistemic modals behave differently from disjunctions under deontic modals, but the burden lies on BD-supporters to present such an account.

### **Procrastinate.**

The canonical resistance strategy is to suppose that:

(4) Procrastinate ought to accept and write.

(5) Procrastinate ought to accept.

are associated with subtly different contexts. More precisely, the idea is to say that the modal base against which (5) is evaluated does not include worlds in which Procrastinate accepts and writes. When we evaluate (4), those worlds must be in the picture. Writers within the Kratzer tradition have not tackled the Procrastinate puzzle directly, but a modal base switch is often proposed to address some related problems. Portner (2009) convincingly makes this move to block Åqvist’s (1967) version of the Good Samaritan Paradox:

(22) It ought to be that John is not injured in a robbery.

[Åqvist’s intended formalization:  $\sim \bigcirc (I)$ ]

(23) It ought to be that someone helps John, who was injured in a robbery.

[Åqvist’s intended formalization:  $\bigcirc (I \ \& \ H)$ ].

Portner points out that Åqvist’s formalization of (23) does not appear to be the best formalization of the sentence: the restrictive relative clause is probably best interpreted as flagging a presupposition, rather than a conjunct in the scope of the

○ operator. If it is a presupposition, then it is naturally incorporated in the modal base against which (23) is interpreted.

In the Procrastinate case, however, the contextual strategy appears *ad hoc*. What principled motivation do we have to assume that the *accept-and-write*-worlds are excluded from the modal base? They certainly are salient possibilities, given Procrastinate's predicament. Of course, if participants to the conversation take it to be impossible that Procrastinate accepts and writes or simply know that he won't, then those worlds should be excluded from the modal base. However, the scenario of Procrastinate did not suggest that accepting and writing was impossible; it merely stated that it was unlikely. Moreover, the relevant intuitions do not require that it be *extremely* unlikely either. One cannot simply exclude unlikely worlds from appearing in modal bases: it would be *ad hoc* and also make (4) false. Neither does the scenario require that participants to the conversation know that Procrastinate won't accept and write.

Some may think that considerations of charity can help pick out the right modal base. However, any such argument needs to be specific about the mechanics of the relevant charity principle, to avoid overgenerating: after all, almost any *ought*-sentence is true against *some* modal base or other. To my knowledge, no such appeals to charity are specific enough. I conclude that resistance strategies that apply neatly to some versions of the Good Samaritan Paradox do not carry over cleanly to the Procrastinate case.

### Conditional 'Oughts'.

Perhaps, a sentence of the form  $\lceil \text{○}p \rceil$  requires that the modal base be compatible with both  $p$  and  $\sim p$  (this is sometimes called a *diversity* condition on the modal base). This requirement could either be part of the content of 'ought' or some kind of presupposition. Be that as it may, the requirement, by itself, fails to solve the problem. Consider:

(24) If either you drink poison or you fail to drink because you are struck by lightning, then you ought to drink poison.

(24) does not seem much better than (7), but it meets the requirement, since the modal base against which the 'ought' is evaluated contains both worlds in which the prejacent is true and worlds in which it is false.

This does not imply that I reject the diversity condition. In fact, assuming some such pragmatic requirement may improve the coverage and the plausibility of my own account. To see this, notice that the explanation I gave earlier, does not clearly apply to cases like,

(25) If you drink a cup of coffee, you ought to drink a cup of coffee.

Considering (24) and (25) it is reasonable to claim that there really are at least two distinct phenomena here. One captured by my theory (the reason why (24) is not guaranteed to be true is that drinking poison does not meet the permissibility

benchmark), the other plausibly captured by the diversity constraint (the reason why (25) does not appear to be acceptable is that, once we restrict to antecedent worlds, there is no alternative to drinking coffee).

The Conditional ‘Oughts’ problem can also be addressed by adopting an account of conditionals that does not satisfy RESTRICTION. The theoretical space here is so broad and RESTRICTION so entrenched, that this option is best left for separate treatment.<sup>31</sup>

## 7 Conclusion.

My contribution to the semantics of ‘ought’ is the development of an account that reconciles COARSENESS and a violation of INHERITANCE in an attractive way. Upholding COARSENESS does not require us to subscribe to some package of disunified resistance strategies in order to answer the problems with INHERITANCE. This sort of account is especially welcome given that the answers to the objections against INHERITANCE do not at present appear convincing.

My proposal draws its inspiration from extant contrastive approaches for ‘ought’. However, I improved upon them by drawing on the technical resources of a ‘modal resolution’ approach. This approach is a key ingredient to my account of COARSENESS, and to my argument that my approach can be viewed as a generalization of the standard approach. Moreover, the modal resolution approach requires less fluctuation in the underlying set of alternatives (even in otherwise unproblematic cases).

The particular proposal I have made is not the only possible implementation of the informal ideas I have presented. Numerous variations suggest themselves (e.g. modeling every parameter in the semantics as determined by a conversational background, importing decision-theoretic machinery in the ranking of options) but their motivation and implementation will be left for further work.

## Notes

\*My biggest thanks are to Niko Kolodny, John MacFarlane and Paolo Santorio who have read and commented on multiple versions of this project over the years. I have also benefitted, at various different stages, from conversation with and comments from Mike Caie, Matthew Chrisman, David Ebrey, Branden Fitelson, Aidan Gray, Mitchell Green, Lloyd Humberstone, Stefan Kaufmann, Chris Kennedy, John Garthoff, Jeff Horty, Peter Ludlow, Andrew Reisner, Daniel Rothschild, Mark Schroeder, Magda Schwager, Justin Snedegar, Daniel Star, James Stazicker, Tim Sundell, Sarah Zobel, a speedy and insightful anonymous referee, as well as audiences at ANU, Barnard College, Berkeley, Boston University, Göttingen (Linguistics), Latrobe University, Northwestern University, the University of Chicago, the 2007 AAP, the 2008 Pacific APA, Semantics and Philosophy in Europe I. Work on this paper was partly funded by NSF project *Fallibility and Revision in Science and Society*, Award No. SES - 0823418, whose support I gratefully acknowledge.

<sup>1</sup> In fact, I think that the explanatory target is even stronger: a complete analysis of ‘ought’ must help explain not only how more specific *ought*-sentences and coarser ones can be jointly felicitously asserted, but also how the former can explain or support the latter. For example, a good theory must explain how the fact that you ought to drive less than 65 mph on this road can (sometimes) be used to support the claim that you ought to drive less than 100 mph. The theory I advance in this paper can be incorporated in an account of this further phenomenon, but working out the details would require us to get clearer on the question of what we mean when we say that an *ought*-sentence explains or supports another *ought*-sentence. For this reason I limit myself to accounting for COARSENESS understood as a purely semantic thesis.

<sup>2</sup>This terminology is a bit abusive because there are many theories that imply INHERITANCE but do not explicate ‘ought’ as a universal quantifier over possible worlds (see, e.g. Lewis, 1974). Insofar as there are problems with INHERITANCE, they also apply to these theories.

<sup>3</sup>Ross (1941). In its original formulation, it was an argument involving *imperative* sentences, rather than declarative *ought*-sentences. The imperative version of the puzzle is discussed more often: two (contrasting) recent treatments are Mastop (2005) and Vranas (ms.). In this paper, I refrain from drawing any conclusions about the imperative version of the puzzle.

<sup>4</sup>Prof. Procrastinate’s case has some structural similarities with the Good Samaritan paradox in deontic logic. See Åqvist (1967) for a discussion, though the literature on the paradox goes back to the work of A.N. Prior.

<sup>5</sup>Frank (1997); Zvolenszky (2002, 2006).

<sup>6</sup>This defect would be minor, if the resistance strategies were genuinely independently motivated. However, as I will show, the resistance strategies require stipulative elements to even get off the ground.

<sup>7</sup>Consider an analogy: Kripke-semantics, the bare formal machinery, does not immediately provide a story about modality in natural language. It does so only when the formal algebraic semantics is connected to an account of truth-at-a-context and some metasemantic theses. The positive contribution of this paper could (but need not) be understood as producing one way of setting up a bridge principle between *neighborhood* semantics and natural language semantics.

<sup>8</sup>The less arcane term ‘complement’ is sometimes used to denote this, but it creates confusions with other parts of linguistics.

<sup>9</sup>Incidentally, it is a mistake to suppose that a semantic framework using orderings of worlds can only make sense to a consequentialist— a mistake effectively dismissed by David Lewis (1978):

The semantic analysis tells us what is true (at a world) under an ordering. It modestly declines to choose the proper ordering. That is work for a moralist, not a semanticist. [...] For instance, a simplistic non-utilitarian might fancy an ordering on which the better of any two worlds is the one with fewer sins. (It is up to him to tell us how he divides the totality of sins into distinct units).

<sup>10</sup>See Jackson and Pargetter (1986) and Finlay (2009) for two such examples. See Snedegar (forthcoming) for an up to date attempt to defend contrastive approaches about a much broader class of deontic modalities.

<sup>11</sup>Some relevant literature on what counts as ‘alternatives’ to a given action dates back to discussions of utilitarianism in the ‘60s and 70s: see, for example, Bergström (1976), Perloff (1979).

<sup>12</sup>In the account I sketched earlier, this would happen by making salient the *negative value* of burning down the philosophy department

<sup>13</sup>This is not to say that conditional antecedents cannot affect the truth-conditions of denotic consequents.

<sup>14</sup>Additionally, I do not think that Jackson and Pargetter’s account provides a good explanation of Ross’s puzzle. Recall that on Jackson’s view  $\lceil \bigcirc p \rceil$  is true iff the  $p$ -world that is closest to the actual world is better than the closest  $o$ -world for each alternative option  $o$ . Suppose that  $p$ ,  $q$  and  $o$  are incompatible options. Suppose also that the closest  $p$ -world is better than the closest  $o$ -world, and that both are better than the closest  $q$ -world. Therefore  $\lceil \bigcirc p \rceil$  is true. Whether or not  $\lceil \bigcirc(p \vee q) \rceil$  is true at this point depends on which is closer among the closest  $p$ -world and the closest  $q$ -world; however, the Ross intuitions do not appear to turn on the relative proximity of these two worlds.

<sup>15</sup>As Lewis (1972) observes, recognizing an important theoretical role for intensions need not imply identifying intensions with contents. All that we assume in the following is that *contents, whatever they may be, determine intensions*.

<sup>16</sup>We can think of this step in two, equally influential, ways. According to the first perspective (cf. Frank, 1997), relativized modals are like generalized quantifiers. A restrictor (which can be explicit or implicit) narrows down the background set of worlds to those worlds that satisfy a certain condition (e.g. worlds that are compatible with a certain type of evidence, or a certain class of desires, and so on). The modal’s effect is to express some quantificational condition over these worlds—so that (12) ends up having truth conditions paraphrasable as: *at every world compatible with the salient group’s evidence, John is at the movies*. This picture represents modals as functions from pairs of propositions to truth-values and relates Kratzer’s framework to the tradition of ‘dyadic deontic logic’. For this tradition, see Lewis (1974, 1981); Chellas (1974); van Fraassen (1972); for a critique of this approach Bonevac (1998). According to the other perspective (which I adopt in the following) a modal has a semantic value relative to a discourse-level parameter  $\mathcal{M}$  and the *background* is implicitly built into the  $\llbracket \cdot \rrbracket^{c,w}$  function.

<sup>17</sup>The background determines  $\mathcal{M}_c$  by *intersection*: if the modal base consists of propositions  $P_1, \dots, P_n$ , the set of epistemically accessible worlds is the intersection of the intensions of those propositions. If propositions just *are intensions*, the set of deontically ideal worlds is  $P_1 \cap \dots \cap P_n$ ; if propositions are not intensions, but determine intensions by some function  $\delta$ , then the set of deontically ideal worlds is  $\delta(P_1) \cap \dots \cap \delta(P_n)$ .

<sup>18</sup>Kolodny and MacFarlane (2010) use these functions directly as primitive components of their semantic theory rather than define them in terms of the ordering. The differences between these two approaches won’t matter here.

Three more complications arise here that do not concern me much, but must be recorded:

- Kratzer takes the ranking to be simply a partial ordering. In this case ‘best’ worlds means ‘maximal’ worlds (i.e. worlds  $w$  such that no other world  $v$  is better than them).
- The ranking may lack maximal worlds as well (e.g. if for every world  $w$ , there is a world that is ranked higher than  $w$ ). See Kratzer (1981) or the presentation of her work in Kaufmann *et al.* (2006) for the canonical solution. See Swanson (forthcoming) for some criticisms of the canonical solution and an alternative approach.
- Typically,  $\mathcal{M}$  and  $\beta$  are relativized to the world of evaluation, so as to allow for contingently true ‘ought’ statements.

<sup>19</sup>A proper treatment of this dependence would require us to introduce *indices* of evaluation and treat modal bases as parameters in the index that can be shifted away from their initial values. In order to keep as simple as possible, I refrained from introducing this complication in the formal accounts in the main text. I invite the reader who feels the force of the problem to ‘translate’ the framework (in this and the next section) in this way.

<sup>20</sup>The direct inspiration for this discussion is Belnap *et al.* (2001), though I do not attribute it to them, because I have extrapolated rather liberally and selectively from their ideas.

<sup>21</sup>Representing actions purely as propositions loses an important piece of information: actions have *agents*, and propositions do not. This choice is not hugely important in this context because the semantic model of this paper built on the simplifying assumption that, in any given context, there is only one agent. In a more general model, there are reasons to represent actions using properties. For an argument to this effect, see Chierchia (1989), for some ideas about the pragmatics of deontic modals that extensively use a framework with properties, see Portner (2007). Schroeder (2010) argues for a syntactic analysis of deliberative *ought*-sentences that would also be best matched by a model on which action types are properties. Even if this is correct, the semantics I propose in this section is flexible enough to be adapted to this alternative syntax.

<sup>22</sup>One can even keep the ordering source as a set of propositions  $\{O_1, \dots, O_j\}$ . Options are ranked according to how many elements of the ordering source they entail. If there is an option that entails every member of  $\{O_1, \dots, O_j\}$ , that option will among the most favored ones.

<sup>23</sup>There are two possible entries for permissibility (which I'll symbolize as  $[\mathcal{P}]'$ ). The first is:

$$\llbracket [\mathcal{P}]_1 p \rrbracket^{c,w} = 1 \text{ iff some } o \geq \text{benchmark is such that } [p]^c \text{ holds at } o.$$

The second:

$$\llbracket [\mathcal{P}]_2 p \rrbracket^{c,w} = 1 \text{ iff some } >\text{-maximal } o \text{ is such that } [p]^c \text{ holds at } o.$$

Absent further constraints, neither account implies the duality of ' $\circ$ ' and  $[\mathcal{P}]'$ . In the next section, I discuss whether this is a problem.

<sup>24</sup>It is not essential to the account that (iii) be part of the lexical entry: it could alternatively be viewed as a presupposition of utterances of *ought*-sentences.

<sup>25</sup>Duality fails on the logics based on the so-called *stit* framework, see Belnap *et al.* (2001), but not quite in ways that block the current argument.

<sup>26</sup>*Proof:* suppose  $p \models q$ , then  $\sim q \models \sim p$ , by contraposition of entailments. By (PI), we have  $[\mathcal{P}] \sim q \models [\mathcal{P}] \sim p$ . Switching to the duals:  $\sim \circ q \models \sim \circ p$ . Contraposing the entailment again we get INHERITANCE.

<sup>27</sup>Both of the accounts I entertained in footnote 23 above entail (PI). On the other hand, as Niko Kolodny points out to me, in the Procrastinate case, we are willing to assert that Procrastinate may accept and write (in the deontic sense of 'may'), but not that he may accept. If that is correct, and if it requires a semantic explanation we may even have room to reject PI itself.

<sup>28</sup>Wedgwood's examples are not actually (2) or (3), but rather Ross's original examples. I doubt that the different examples matter to the substantial points. Wedgwood's attitude to Ross's puzzle is widespread and anticipated in Føllesdal and Hilpinen (1971).

<sup>29</sup>Simons (2007) is a fully developed BD-account.

<sup>30</sup>Note that free choice effects arise indifferently for epistemic and deontic interpretations of the modals, so a story about free-choice interpretations does not seem to suffice for an account of Ross's puzzle.

<sup>31</sup>The conventional approach (Frank, 1997) here is to propose that in addition to an overt deontic modal 'If ... ought ...' conditionals also have a covert epistemic necessity modal, just 'above' the overt deontic operator. The conditional antecedent restricts the higher, but not the lower operator, whose modal base is allowed to reach outside of the antecedent worlds. A full development of this idea requires representing Kratzer's approach in a somewhat deeper way than I have done here (in particular, it requires having modal bases be not just sets fixed directly by context, but have them vary with the world of evaluation). I am skeptical of this solution: I do not find its independent motivation convincing and I am convinced by the criticisms in Zvolenszky (2002): it can only succeed at the price of overgenerating the readings of other 'If ... ought...' conditionals and making predictions. A fully general resolution of that issue is beyond the scope of this paper (and Zvolenszky's). See the unpublished Geurtz (unpublished), for an opposing voice in this debate.



## Bibliography

- ÅQVIST, Lennart (1967), 'Good Samaritans, Contrary-to-duty Imperatives, and Epistemic Obligations', *Noûs*, 1(4), 361–379.
- BELNAP, Nuel, PERLOFF, Michael, and XU, Ming (2001), *Facing the Future* (Oxford University Press).
- BERGSTRÖM, Lars (1976), 'On the Formulation and Application of Utilitarianism', *Noûs*, 10(2), 121–144.
- BONEVAC, Daniel (1998), 'Against Conditional Obligation', *Noûs*, 32(1), 37–53.
- CHELLAS, Brian (1974), 'Conditional Obligation', in Soren Stenlund (ed.), *Logical Theory and Semantic Analysis: Essays Dedicated to Stig Kanger on His Fiftieth Birthday* (Reidel).
- (1980), *Modal Logic* (Cambridge University Press).
- CHIERCHIA, Gennaro (1989), 'Anaphora and Attitudes De Se', in R. Bartsch, J. van Benthem, and P. van Emde Boas (eds.), *Semantics and Contextual Expression* (Groeningen-Amsterdam Studies in Semantics).
- DUMMETT, Michael (1973), 'The Justification of Deduction', in *Truth and Other Enigmas (1981)* (Duckworth).
- FINLAY, Stephen (2009), 'Oughts and Ends', *Philosophical Studies*, 143(3), 315–340.
- FÖLLESDAL, Dagfin and HILPINEN, Risto (1971), 'Deontic Logic: an Introduction', in R. Hilpinen (ed.), *Deontic Logic: Introductory and Systematic Readings*, 1–35 (Reidel).
- FRANK, Annette (1997), *Context Dependence in Modal Constructions*, Ph.D. thesis, University of Stuttgart.
- GEURTZ, Bart (unpublished), 'On an Ambiguity in Quantified Conditionals', University of Nijmegen.
- JACKSON, Frank (1985), 'On the Semantics and Logic of Obligation', *Mind*, 94(374), 177–195.
- JACKSON, Frank and PARGETTER, Robert (1986), 'Oughts, Options and Actualism', *The Philosophical Review*, 95(2), 233–255.
- KAPLAN, David (1989), 'Demonstratives', in J. Almog, J. Perry, and H. Wettstein (eds.), *Themes from Kaplan* (Oxford University Press).
- KAUFMANN, Stefan, CONDORAVDI, Cleo, and HARIZANOV, Valentina (2006), 'Formal Approaches to Modality', in W. Frawley (ed.), *The Expression of Modality* (Mouton-de Gruyter).

- KOLODNY, Niko and MACFARLANE, John (2010), 'Ifs and oughts', *Journal of philosophy*, 107(3), 115–143.
- KRATZER, Angelika (1977), 'What 'Must' and 'Can' Must and Can Mean', *Linguistics and Philosophy*, 1(3), 337–355.
- (1981), 'The Notional Category of Modality', in B. Partee and P. Portner (eds.), *Formal Semantics: the Essential Readings* (Blackwell).
- (1991a), 'Conditionals', in A. von Stechow & D. Wunderlich (ed.), *Semantics: An International Handbook of Contemporary Research* (De Gruyter). From the *Semantics Archive*.
- (1991b), 'Modality', in A. von Stechow & D. Wunderlich (ed.), *Semantics: An International Handbook of Contemporary Research* (De Gruyter).
- LEWIS, David K. (1972), 'General Semantics', in D. Davidson and G. Harman (eds.), *Semantics for Natural Language* (Reidel).
- (1974), 'Semantic Analyses for Dyadic Deontic Logic', in *Papers in Ethics and Social Philosophy* (Cambridge University Press).
- (1978), 'Reply to McMichael', *Analysis*, 38(2), 85–86.
- (1981), *Counterfactuals* (Blackwell).
- (1996), 'Elusive Knowledge', *Australasian Journal of Philosophy*, 74(4), 549–567.
- MASTOP, Rosja (2005), *What Can You Do?*, Ph.D. thesis, ILLC Amsterdam.
- PERLOFF, Michael (1979), 'A Better Alternative', *Analysis*, 39(2), 106–108.
- PORTNER, Paul (2007), 'Imperatives and Modals', *Natural Language Semantics*, 15(4), 351–383. Georgetown University.
- (2009), *Modality* (Oxford University Press).
- ROSS, Alf (1941), 'Imperatives and Logic', *Theoria*, 7, 53–71.
- SCHROEDER, Mark (2010), 'Oughts, Agents and Actions', *Philosophical Review*, 120(1), 1–41.
- SIMONS, Mandy (2007), 'Semantics and Pragmatics in the Interpretation of *or*', *Proceedings of SALT 15*.
- SLOMAN, Aaron (1970), 'Ought and Better', *Mind*, 79(315), 385–394.
- SNEDEGAR, Justin (forthcoming), 'Contrastive Semantics for Deontic Modals', *forthcoming in Contrastivism in Philosophy*, M. Blaauw (ed.), Routledge.

- SWANSON, Eric (forthcoming), ‘On the Treatment of Incomparability in Ordering Semantics’, *Journal of Philosophical Logic*.
- VAN FRAASSEN, Bas (1972), ‘The Logic of Conditional Obligation’, *Journal of Philosophical Logic*, 417–438.
- (1973), ‘Values and the Heart’s Command’, *Journal of Philosophy*, 70(1), 5–19.
- VRANAS, Peter (ms.), ‘New Foundations for Imperative Logic II: Pure Imperative Inference’, University of Wisconsin, Madison.
- WEDGWOOD, Ralph (2006), ‘The Meaning of ‘Ought’’, *Oxford Studies in Metaethics*, 1, 127–160.
- YALCIN, Seth (2009), ‘Nonfactualism About Epistemic Modality’, in Egan A. and Weatherson B. (eds.), *Epistemic Modality* (Oxford University Press).
- ZVOLENSZKY, Zófia (2002), ‘Is a possible worlds semantics of modality possible? A problem for Kratzer’s semantics’, in B. Jackson (ed.), *Proceedings of SALT 12*, 339–358.
- (2006), ‘A Semantic Constraint on the Logic of Modal Conditionals’, in Beáta Gyuris, László Kálmán, Chris Pi non, and Károly Varasdi (eds.), *Proceedings of the Ninth Symposium on Logic and Language (LoLa 9)* (Révai Digital Press).