

Research Bank

Journal article

Inheritance : Professor Procrastinate and the logic

Blumberg, Kyle and Hawthorne, John

This is the peer reviewed version of the following article:

Blumberg, K. and Hawthorne, J. (2023). Inheritance : Professor Procrastinate and the logic. *Philosophy and Phenomenological Research*. 106(1), pp. 84-106, which has been published in final form at <https://doi.org/10.1111/phpr.12846>.

This article may be used for non-commercial purposes in accordance with Wiley Terms and Conditions for Use of Self-Archived Versions. This article may not be enhanced, enriched or otherwise transformed into a derivative work, without express permission from Wiley or by statutory rights under applicable legislation. Copyright notices must not be removed, obscured or modified. The article must be linked to Wiley's version of record on Wiley Online Library and any embedding, framing or otherwise making available the article or pages thereof by third parties from platforms, services and websites other than Wiley Online Library must be prohibited.

INHERITANCE: PROFESSOR PROCRASTINATE AND THE LOGIC OF OBLIGATION*

Kyle Blumberg
Dianoia Institute of Philosophy, Australian Catholic University
John Hawthorne
School of Philosophy, University of Southern California
Dianoia Institute of Philosophy, Australian Catholic University

Abstract

Inheritance is the principle that deontic ‘ought’ is closed under entailment. This paper is about a tension that arises in connection with *Inheritance*. More specifically, it is about two observations that pull in opposite directions. One of them raises questions about the validity of *Inheritance*, while the other appears to provide strong support for it. We argue that existing approaches to deontic modals fail to provide us with an adequate resolution of this tension. In response, we develop a positive analysis, and show that this proposal provides a satisfying account of our intuitions.

1 Introduction

Inheritance is the principle that deontic ‘ought’ (and its cognates, e.g. ‘must’, ‘has to’, etc.) is closed under entailment:¹

Inheritance *If* $p \models q$, *then* $\text{ought } p \models \text{ought } q$

This paper is about a tension that arises in connection with *Inheritance*. More specifically, it is about two observations that pull in opposite directions. One of them raises questions about the validity of *Inheritance*, while

*To appear in *Philosophy and Phenomenological Research*.

¹We assume that ‘ought’ takes a sentential, proposition-denoting complement. This is common practice in philosophical semantics, but it isn’t essential to our arguments. For instance, if one thinks that ‘ought’ instead takes predicates (Schroeder, 2011), then one could formulate *Inheritance* using a notion of predicate entailment (Schlenker, 2009). The rest of the dialectic would then proceed as below, *mutatis mutandis*.

the other appears to provide strong support for it. To illustrate the first, consider a famous case originally discussed by Jackson & Pargetter (1986), and made sharper by Cariani (2013):

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 (but not certain) 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.

- (1) a. Prof. Procrastinate ought to accept.
- b. Prof. Procrastinate ought to accept and write.

Many have the intuition that (1a) sounds false in context. After all, if Procrastinate were to accept, he’d very likely not write the review, which would be very bad. By contrast, (1b) is clearly true, since Procrastinate accepting and writing is the best outcome. But this constellation of judgments puts pressure on Inheritance, since *Procrastinate accepts and writes* obviously entails *Procrastinate accepts*.

On the other hand, theorists have also pointed out that one cannot felicitously conjoin (1b) with the negation of (1a) (Crnič, 2011; von Fintel, 2012):²

²We use ‘have to’ in these examples since English constructions which feature negation outscoping ‘ought’ tend to sound awkward, e.g. ‘It’s not the case that Procrastinate ought to accept and write’ (von Fintel, 2012; Iatridou & Zeijlstra, 2013). Another way to get around this problem is to use verbs that lexicalize negation, e.g. ‘doubt’, as in (2):

- (2) a. # I’m certain Procrastinate ought to accept and write, but I doubt he ought to accept.
- b. # I doubt Procrastinate ought to accept, but I’m certain he ought to accept and write.

But it is more straightforward to work with unembedded sentences, so we stick with the examples in (3). Note that there are well-known differences between ‘ought’ and ‘have to’ (von Fintel & Iatridou, 2008). However, these aren’t relevant for our purposes, so we ignore them here.

We also note that analogues of (3) where negation takes narrow-scope in the second conjunct are unacceptable, for example ‘Prof. Procrastinate ought to accept and write, but he ought not accept’. These examples also militate in favor of Inheritance given the plausible principle of “deontic non-contradiction”: $\lceil \text{ought } p \rceil$ and $\lceil \text{ought } \neg p \rceil$ can never both be true.

- (3) a. # Prof. Procrastinate has to accept and write, but he doesn't have to accept.
b. # Prof. Procrastinate doesn't have to accept, but he has to accept and write.

Both (3a) and (3b) are unacceptable (as indicated by the ‘#’ preceding each example). But this is unexpected if Inheritance fails: if (1a) is false and (1b) is true, why can't one conjoin them as in (3)? Indeed, the unacceptability of these conjunctions appears to provide strong abductive support for Inheritance. If Inheritance is valid then we have a straightforward explanation for why these sentences sound bad: they can never be true.³

We argue that existing approaches to deontic modals fail to provide us with a satisfactory resolution of the tension exhibited by (1) and (3). In short, accounts which reject Inheritance have a hard time explaining why the conjunctions in (3) are bad. And existing theories which validate Inheritance don't yield an adequate explanation for why (1a) sounds false in context.

In response, we develop a positive analysis of deontics. We will try to show that our proposal does a better job of accounting for the intuitions elicited by (1) and (3). In broad terms, our semantics validates Inheritance, and therefore bears an affinity to standard, quantificational approaches. Consequently, we are able to provide a straightforward explanation for why the conjunctions in (3) are infelicitous. However, on our account Inheritance is not valid *tout court*, or classically valid. Instead, it is only valid if we assume that certain background conditions, or presuppositions are satisfied. And as we will show, it is precisely the content and character of these presuppositions which allow us to explain the contrast between (1a) and (1b).

The paper is structured as follows. In §2 we consider existing approaches to deontics, and argue that they fail to provide an adequate explanation of our observations. Then in §3 we develop our positive proposal in several stages. Finally, §4 concludes.

2 Existing Accounts

In this section, we consider some of the most prominent existing analyses of deontic modals. More specifically, in §2.1 we present two accounts on which Inheritance is invalid, and in §2.2 we consider a popular theory which

³Some authors distinguish between “objective” and “subjective” senses of deontics such as ‘ought’, e.g. Prichard (1932); Ross (1939); Jackson (1991, 1997). However exactly one understands this distinction, we submit that analogues of (2) and (3) can be constructed for each sense. So, our puzzle involving Inheritance doesn't hang on any particular interpretation of ‘ought’.

validates Inheritance. We show that none of these accounts successfully resolves the tension discussed in §1.

2.1 Actualism and decision-theoretic accounts

The most popular Inheritance-rejecting accounts are types of *contrastivism* about deontic claims. Contrastivists maintain that sentences such as (1a) (‘Prof. Procrastinate ought to accept’) have an implicitly comparative nature. More precisely, contrastivism is the idea that (1a) is true just when the proposition *Procrastinate accepts* is better than the relevant alternative propositions. A central point of difference between contrastivists concerns how the notion of “betterness” should be understood, i.e. what grounds the comparison between propositions. We will briefly consider two varieties of contrastivism: actualism and decision-theoretic accounts.⁴

Actualists spell out the comparison in terms of similarity to the actual world (or more generally, the world of evaluation).⁵ They also take for granted a preference ordering over possible worlds, so that it makes sense to talk about one world being better than another. If we suppose that the only alternative relevant for the evaluation of ‘ought p ’ is $\neg p$, actualism can be presented as follows:⁶

Actualist semantics for *ought*

‘ought p ’ is true in w iff the closest p -world to w is better than the closest $\neg p$ -world to w

By contrast, decision-theoretic accounts spell out the comparison in terms of decision-theoretic notions such as *expected value*.⁷ The expected value of p is the utility of p weighted by the relevant probabilities (Jeffrey, 1965).⁸ Again, assuming that the only alternative relevant to the evaluation of ‘ought p ’ is $\neg p$, a simple decision-theoretic account can be presented as follows:

Decision-theoretic semantics for *ought*

‘ought p ’ is true in w iff the expected value of p in w is higher than the expected value of $\neg p$ in w

⁴Also see (Finlay, 2019) for a further contrastivist theory that rejects Inheritance.

⁵See, e.g. (Jackson, 1985; Jackson & Pargetter, 1986).

⁶We let ‘ p ’ range over the logical forms of proposition-denoting strings and let ‘ w ’ range over the corresponding propositions denoted by ‘ p ’.

⁷See e.g. (Goble, 1996; Lassiter, 2017).

⁸Note that if there are many ways for an agent to act so that p obtains, then the underlying probability distribution will have to assign probabilities to the acts themselves.

It is fairly easy to see that both actualism and decision-theoretic accounts can explain our first observation from §1. This concerned the fact that (1a) sounds false, but (1b) sounds true in the *Procrastinate* scenario:

- (1) a. Prof. Procrastinate ought to accept.
- b. Prof. Procrastinate ought to accept and write.

Assuming that the closest world to actuality where Procrastinate accepts is one where he fails to live up to his obligations, actualists predict that (1a) should be false. By contrast, the closest world where Procrastinate accepts *and* writes will be better than the closest world where he fails to accept and write. So, actualists predict that (1b) should be true. Similarly, since Procrastinate will likely not complete the review if he accepts, the expected value of him accepting is very low. So, the decision-theoretic account predicts that (1a) should be false. On the other hand, the expected value of Procrastinate accepting *and* writing is very high, so the decision-theoretic account predicts that (1b) should be true.

However, as von Fintel (2012) points out, theories such as actualism and the decision-theoretic account leave our second observation unexplained. If (1a) is false and (1b) is true, it is completely mysterious why conjunctions such as (3a) should sound so terrible:

- (3a) # Prof. Procrastinate has to accept and write, but he doesn't have to accept.

The problem is quite general: it is difficult to see how *any* account that makes (1a) false and (1b) true (in a single context) could explain why the target conjunctions are abominable.^{9,10} This suggests to us that there is

⁹Actualists might argue that (3a) is predicted to be false on their account so long as the improbable happens, and Procrastinate ends up accepting and writing. In this case, (1a) will be true, assuming that the actual world is closest to itself. But in this scenario the actualist also predicts that unacceptable speeches such as 'Although we didn't realize it back then, it was true all along that Procrastinate ought to have accepted' should be true. Moreover, even if we know that Procrastinate won't both accept and write (as in Jackson & Pargetter's original case), (3a) is infelicitous. However, granted plausible assumptions, in this context the actualist predicts that (1a) is false and (1b) is true.

¹⁰In passing, von Fintel (2012) suggests that the Inheritance-denier might try to appeal to (Moss, 2012) in order to explain what is going on with examples such as (3a). But there is an immediate difficulty with this idea. Moss's account is aimed at trying to explain order effects arising from counterfactual conditionals, e.g. why (4) sounds fine but (5) is infelicitous:

- (4) If Mary goes to the concert, she'll have fun, but if she goes to the concert and is stuck behind a tall person, she won't have fun.
- (5) # If Mary goes to the concert and is stuck behind a tall person, she won't have

something not quite right with these approaches to deontic modals.¹¹

Although we find actualism and decision-theoretic accounts problematic, it is worth mentioning that we still think there is something fundamentally correct about the basic intuition underlying contrastivism, namely that deontic claims are in some sense comparative by nature. But as we will see in §3, our preferred way of implementing this idea is compatible with Inheritance, and therefore predicts that examples such as (3a) should be unacceptable.¹²

2.2 Kratzer's account

Accounts that validate Inheritance are by and large versions of the standard, Kratzerian analysis of modal expressions in natural language. On this approach, modals operate as quantifiers over a usually restricted domain of possibilities. This domain is fixed by two parameters: (i) the *modal base*, which determines the broad set of possibilities relevant for evaluating the modal, and (ii) the *ordering source*, which provides an ordering of the possibilities in the modal base. On this view, different modal flavors (e.g. epistemic, deliberative, deontic, bouletic, etc.) correspond to different values of the two parameters.¹³

Kratzer's approach makes use of a function $\text{BEST}(\cdot, \cdot)$ that takes a modal base and an ordering and yields the set of top-ranked worlds as determined by the ordering. We can simplify by restricting our attention to scenarios where the modal base is finite. In this case, BEST can be spelled out as follows:¹⁴

Specification of BEST

For any modal base \mathcal{B} , and ordering $>$:

$$\text{BEST}(\mathcal{B}, >) = \{w' \in \mathcal{B} \mid \neg \exists w'' \in \mathcal{B} \text{ such that } w'' > w'\}$$

fun, but if she goes to the concert, she'll have fun.

However, the abominable conjunctions that we are interested in don't exhibit an order effect; (3b) ('Prof. Procrastinate doesn't have to accept, but he has to accept and write) sounds just as bad as (3a).

¹¹Even Cariani (2021, 27), who himself rejects Inheritance, admits that 'little has been said by critics of Inheritance, myself included, to account for [examples such as (3a)] in a systematic way'.

¹²For further criticisms of actualism and decision-theoretic accounts, see (Cariani, 2016a).

¹³See, e.g. (Kratzer, 1977, 1981, 1991, 2012).

¹⁴We only focus on finite cases for simplicity. There are well-known ways of formulating BEST in order to cover cases where the ordering fails to contain maximal elements (Kratzer, 2012).

Then Kratzer’s account of deontics is essentially the following:¹⁵

Kratzer’s semantics for *ought*

‘ought p ’ is true relative to $\langle w, \mathcal{B}, < \rangle$ iff every $w \in \text{BEST}(\mathcal{B}, <)$ is a p -world

This analysis makes Inheritance valid: if all of the best worlds are p -worlds, and p entails q , then all the best worlds will be q -worlds as well. So, no matter what the modal base and ordering is, the conjunctions in (3) can never be true. This provides a straightforward explanation for their infelicity.

But as it stands, the account can’t explain our first observation from §1. Any context in which (1b) (‘Prof. Procrastinate ought to accept and write’) is true will also be one in which (1a) (‘Prof. Procrastinate ought to accept’) is true as well. Why, then, is the latter infelicitous? When discussing similar cases, proponents of the Kratzerian analysis generally appeal to a context shift here; they maintain that (1a) and (1b) are evaluated relative to distinct parameters. For instance, von Fintel (2012) claims that the modal base shifts. (1b) is evaluated relative to a modal base that includes worlds where Procrastinate writes and accepts. Since such worlds are best, the sentence is true. By contrast, (1a) is evaluated relative to a modal base that excludes the possibilities in which Procrastinate accepts and writes, these worlds are ignored in context. The best worlds in this restricted set are ones where the professor declines, so (1a) is false.

Bronfman & Dowell (2018) also invoke context dependence, but locate it in a different place. They maintain that (1a) and (1b) are evaluated relative to distinct orderings. (1b) is evaluated relative to an ordering which only cares about the overall goodness of the outcome, and thus ranks worlds in which Procrastinate both accepts and writes as top. Relative to this ordering, (1b) is true. On the other hand, (1a) is evaluated relative to an ordering on which ‘a world w gets ranked on the basis of how well Procrastinate can be expected to discharge his professional responsibilities, given the action he performs in w during the time it would take him to accept’ (p.96). Relative to this ordering, the best worlds are ones where Procrastinate declines, so (1a) is false.

One might worry that these responses in defense of the Kratzerian analysis undermine the account’s explanation of what goes wrong with abominable conjunctions. If the parameters relative to which (1a) and (1b) are evaluated

¹⁵Strictly speaking, on Kratzer’s account deontic claims are evaluated relative to more complex objects than just a set of worlds and an ordering. Working with these more complex elements is important for some applications—see, e.g. (Cariani *et al.*, 2013; Cariani, 2016b). However, they aren’t necessary for our purposes, so we use a simpler representation here.

can shift, then why can't both conjuncts in (3a) ('Prof. Procrastinate has to accept and write, but he doesn't have to accept') be true? Here the Kratzerian can appeal to an independently motivated fact about context-sensitivity: intersentential context shifts are often difficult to achieve. For instance, consider "weather predicates" such as 'raining'. The meaning of these expressions can be modulated by explicit linguistic material, e.g. 'It is raining in London on June 1st, 1994'. But they can also be modulated by contextually salient information, e.g. when I ask my friend in London 'Is it raining?', I will naturally be understood as asking 'Is it raining *in London*?', even if I am in New York City at the time of the call. However, even if both my location and my friend's have been made salient, it is very difficult to interpret each occurrence of 'raining' in (6a) relative to distinct locations. The sentence simply sounds incoherent. Indeed, even if we make one of the locations explicit, as in (6b), it is still hard to recover a good interpretation of the sentence.

- (6) a. # It's raining, but it's not raining.
 b. # It's raining in London, but it's not raining.

Similar observations can be made about other context-sensitive expressions, for example quantifiers and gradable adjectives such as 'tall':

- (7) a. # Everyone is asleep but someone is awake.
 b. # Mark is tall but he is short.

It is usually assumed that the domains of quantifiers are fixed by context (Stanley & Szabó, 2000), and that gradable adjectives are interpreted relative to a contextually determined scale, e.g. a scale of tallness (Kennedy, 2007). However, it is plausible that neither of these parameters are able to shift in the examples in (7), which accounts for why they are infelicitous. Indeed, there is independent evidence that the parameters relative to which modals are interpreted cannot easily be shifted mid-sentence. For example, consider epistemic modals such as 'might' and 'must':

- (8) # Amy must be at work but she might be at home.

Presumably (8) is unacceptable because 'might' and 'must' are being evaluated relative to the same modal base; and relative to a single modal base both conjuncts cannot be true together. Overall, there is compelling evidence that mid-sentence context shifts are often difficult to achieve. So, it is plausible that intersentential context shifts are difficult to bring about in constructions such as (3a). If that's correct, then appealing to a context shift

in response to our observations in (1) doesn't undermine the Kratzerian's account of abominable conjunctions.¹⁶

There are, however, problems with these appeals to context shifting. Let us consider von Fintel's response first. The *Procrastinate* story specifically emphasizes that it is not *impossible* Procrastinate both accepts and writes, just that this outcome is *unlikely*. But on the Kratzerian approach, the modal base for 'ought' is taken to be *circumstantial*—that is, determined by the relevant facts of the situation. There is some vagueness in what exactly “relevant” means here, and how precisely circumstantial modal bases should be specified. However, Procrastinate's dispositions to accept and write seem relevant on almost any plausible way of understanding this notion. So, by the Kratzerian's own lights, worlds in which Procrastinate accepts and writes should appear in the modal base.

In response, it might be argued that low-probability outcomes are often ignored for the purposes of semantic evaluation, and that this is what happens when we assess (1a).¹⁷ But there being a low probability that Procrastinate both accepts and writes isn't necessary to get the relevant judgments going. So long as the outcome where Procrastinate accepts and fails to write is made bad enough, we can make it quite *likely* that Procrastinate accepts and writes. Consider the following variant of the *Procrastinate* scenario:

Procrastinate 2: Prof. Procrastinate is invited to review a book on which he is the only fully qualified specialist on the planet. Procrastinate is generally reliable, but sometimes fails to bring

¹⁶To be clear, intersentential context shifts are not *always* difficult to achieve. As mentioned, it is usually assumed that the domains of quantifiers are determined by context. In that case, examples such as Soames's (1986) (9) involve a fairly seamless mid-sentence context shift, since the domain for the quantifier 'everyone' obviously can't be the same as the domain for the indefinite description:

(9) Everyone is asleep and is being monitored by a research assistant.

This raises the following question: why are intersentential contexts shifts easy to achieve in some cases, e.g. (9), but difficult to accomplish in others, e.g. (6a)-(8)? We are not aware of a systematic theory that addresses this question. But given that there is independent evidence that intersentential context shifts are hard to accomplish in many cases, we don't think it is particularly problematic to maintain that our target abominable conjunctions are cases where each conjunct is, by default, interpreted relative to the same contextual parameter. Indeed, it is especially striking that mid-sentence context shifts are difficult to achieve with other modal expressions, e.g. epistemic modals. This provides at least *prima facie* support for the claim that the contextually determined parameters relative to which deontic modals are evaluated can't be easily shifted either.

¹⁷It is plausible that something like this occurs when we evaluate conditionals. For instance, the truth of 'If I drop this plate, it'll smash' doesn't seem to be undermined by worlds where some quantum event occurs so that the plate lands safely on the floor without breaking (Lewis, 2016).

projects to completion. In particular, if Procrastinate accepts to review the book, it is quite likely (but not certain) that he will end up writing the review. If Procrastinate accepts but doesn't end up writing the review, the consequences will be catastrophic: his actions will precipitate a crisis in tertiary education that will culminate in several universities closing and thousands of academics losing their livelihoods. If Procrastinate declines, someone less qualified than him will end up writing the review, but there will be no crisis.

To our ears, the relevant pattern of judgments in *Procrastinate* is replicated in *Procrastinate 2*. (1a) still sounds false, since the costs of Procrastinate accepting and failing to write the review are enormous. On the other hand, (1b) still sounds true. If that's right, then ignoring the outcome where Procrastinate accepts and writes on the grounds that it is low probability isn't an adequate response to the phenomenon.

Turning to Bronfman & Dowell's reply, recall that they do not appeal to a shift in the modal base. Instead, they appeal to a shift in the ordering that is relevant in context. The claim is that (1a) is false on the ordering relative to which it is evaluated, since all of the best worlds are ones where Procrastinate declines. But this makes bad predictions. It implies that sentences such as (10a) and (10b) should have true interpretations in context:

- (10) a. Prof. Procrastinate ought to/has to not accept and write.
 b. Prof. Procrastinate ought to/has to do something other than accept and write.

However, it is very difficult to hear these as true in the *Procrastinate* scenario.^{18,19}

¹⁸The *Procrastinate 2* scenario also poses a problem for Bronfman & Dowell. Recall that on their account the ordering relative to which (1a) is evaluated is one where 'a world w gets ranked on the basis of how well Procrastinate can be expected to discharge his professional responsibilities, given the action he performs in w during the time it would take him to accept'. However, in *Procrastinate 2* Procrastinate *can* be expected to discharge his professional responsibilities, given that he accepts. So, on both of the orderings that Bronfman & Dowell appeal to, the best worlds are ones where Procrastinate accepts and writes. That said, we agree with Bronfman & Dowell that the unacceptability of (1a) is partly determined by certain temporal considerations—see §3.3 for further discussion.

¹⁹Von Fintel also maintains that any context where (1a) is unacceptable is one where all of the best worlds in the relevant modal base are *decline*-worlds. However, on his approach every world in this modal base is one where Procrastinate fails to both accept and write. This is relevant because proponents of the Kratzerian analysis often impose the following constraint on certain modal claims: $\Box p$ is true with respect to a modal base \mathcal{B} only if \mathcal{B} contains both p -worlds and $\neg p$ -worlds, i.e. p is *diverse* with respect to \mathcal{B} (Condoravdi, 2002). On von Fintel's approach, the examples in (10) fail the diversity

We should say that although von Fintel and Bronfman & Dowell’s attempt to explain the badness of (1a) (‘Prof. Procrastinate ought to accept’) is problematic, we still think that appealing to a context-shift here is essentially correct. As we will see, what distinguishes our approach from the previously mentioned authors is the parameter associated with the shift.

To sum up, existing theories of deontic modals have competing strengths: views that reject Inheritance can explain why (1a) sounds false and (1b) (‘Prof. Procrastinate ought to accept and write’) sounds true, but they can’t explain why conjunctions such as (3a) (‘Prof. Procrastinate has to accept and write, but he doesn’t have to accept’) are abominable. On the other hand, views that validate Inheritance can explain why the relevant conjunctions are abominable, but they can’t provide a satisfying explanation for why (1a) sounds false. In the next section, we develop a positive proposal which, we argue, yields a more promising account of our observations.

3 A semantics for *ought*

We develop our theory in several stages. First, we present our basic entry for ‘ought’ in §3.1, which allows us to explain why conjunctions such as (3a) are unacceptable. Then in §§3.2-3.3 we add presuppositions to our account, and discuss some important meta-semantic considerations. Our discussion in these subsections allows us to explain the contrast between (1a) and (1b). Finally, in §3.4 we contrast our approach with those of Cariani (2013, 2016b).

3.1 The basic proposal

Like Kratzer’s account considered in §2.2, we also propose that ‘ought’ operates as a quantifier over a restricted domain of possibilities. However, one of the key features of our account is that we take the objects in this domain to be *propositions* rather than *worlds*. There are several ways of developing this idea, but we will implement it in a fairly simple way so as not to distract from our central arguments.

We will say that \mathcal{A} is a set of *alternatives* if it is a set of pairwise incompatible propositions. So, if $A, B \in \mathcal{A}$, then $A \cap B = \emptyset$. To illustrate, let ANN, MARY, PETE, and SUE represent the propositions that Ann wins the race, Mary wins the race, Pete wins the race, and Sue wins the race, respectively. Then $\mathcal{A}_1 = \{\text{ANN}, \text{MARY}, \text{PETE}, \text{SUE}\}$ is a set of alternatives. We propose that the

condition relative to the modal base where (1a) is unacceptable. On the other hand, it is worth noting that positing a diversity condition for deontic modals is fairly controversial (Frank, 1997; Zvolensky, 2002).

set of objects that is relevant for the evaluation of an ‘ought’ claim is a set of contextually supplied alternatives.

We also propose that ‘ought’ claims are evaluated relative to a contextually determined ordering \succ over propositions. The function $\text{TOP}(\cdot, \cdot)$ takes a set of alternatives \mathcal{A} and ordering \succ , and returns the maximal elements in the ordering.²⁰ For instance, suppose \succ orders \mathcal{A}_1 as follows:

ANN \succ MARY \succ PETE \succ SUE

Then $\text{TOP}(\mathcal{A}_1, \succ) = \text{ANN}$.²¹

Our first-run account can be expressed as follows:

Account 1

‘ought p ’ is true relative to $\langle w, \mathcal{A}, \succ \rangle$ iff for every $q \in \text{TOP}(\mathcal{A}, \succ)$:
 $q \subseteq p$

In short, ‘ought p ’ is true just in case all of the top-ranked alternatives in the ordering entail p . To be clear, we aren’t the first to propose an alternative-sensitive entry along the lines of Account 1 (Cariani, 2013; Cariani *et al.*, 2013; Cariani, 2016b). But we will develop this approach in a novel direction.²²

To get a feel for how Account 1 works, consider the following famous example from Parfit (1988):²³

Miners: Ten miners are trapped either in shaft A or in shaft B, but we do not know which. Flood waters threaten to flood the shafts. We have enough sandbags to block one shaft, but not both. If we block one shaft, all the water will go into the other shaft, killing any miners inside it. If we block neither shaft, both shafts will fill halfway with water, and just one miner, the lowest in the shaft, will be killed.

- (11) We ought to block neither shaft.

(11) has a salient true reading here, but this can’t be captured by the standard Kratzerian analysis (Cariani *et al.*, 2013). This is because blocking

²⁰For simplicity we assume that \mathcal{A} has finite cardinality.

²¹As a shorthand, we will write $\text{TOP}(\mathcal{A}_1, \succ) = \text{ANN}$ when we mean $\text{TOP}(\mathcal{A}_1, \succ) = \{\text{ANN}\}$.

²²See §3.4 for a comparison between our approach and (Cariani, 2013, 2016b).

²³Parfit credits Regan (1980) with the case. Also see (Kolodny & MacFarlane, 2010) for discussion.

neither shaft is guaranteed to lead to a suboptimal outcome; the best worlds are ones where we block shaft A and the miners are in A, or we block shaft B and the miners are in B. However, alternatives are relatively coarse-grained entities. So, even if there are some *worlds* in an alternative B that are suboptimal, B can still be ranked higher than the other alternatives. For instance, let us suppose that the relevant set of alternatives in the *Miners* scenario is $\mathcal{A} = \{\text{NEITHER}, A, B\}$, where **NEITHER** is the proposition that we block neither shaft, **A** is the proposition that we block shaft A, and **B** is the proposition that we block shaft B. Moreover, let us suppose that the relevant ordering over these alternatives looks as follows:

$$\text{NEITHER} \succ A/B$$

In this case, Account 1 predicts that (11) should be true.

At this point, two natural meta-semantic questions arise for Account 1: (i) how exactly does the set of alternatives \mathcal{A} get determined in context, and (ii) how is the ordering over alternatives \succ structured. We will have more to say about (i) in §3.3. But to a first approximation, the alternatives relevant for the evaluation of an ‘ought’ claim are tied to the actions that the relevant agent(s) can perform, their *decision problem*. Importantly, decision problems become more fine-grained over time, as the space of possible actions available to the agent increases. For instance, in the *Procrastinate* scenario, Prof. Procrastinate’s immediate options involve deciding between accepting and declining. Let us call the set of alternatives corresponding to this decision problem **coarse**, so **coarse** = {ACCEPT, DECLINE}. But relative to a longer time-frame, more actions become available to Procrastinate. In particular, he has two ways of acting once he’s accepted: he can either accept and write, or accept and not write. Let us call the set of alternatives corresponding to this more fine-grained decision problem **fine**, so **fine** = {ACCEPT&WRITE, ACCEPT& $\overline{\text{WRITE}}$, DECLINE}. It is our general contention that there are two relevant sets of alternatives in the *Procrastinate* scenario: **coarse** and **fine**.

As for (ii), we’re attracted to the idea that in the relevant contexts, e.g. *Procrastinate*, the ordering over alternatives is determined by decision-theoretic considerations, e.g. expected value. In that case, for alternatives A, B : $A \succ B$ when the expected value of A is greater than the expected value of B . This would explain why, for example, **NEITHER** is top-ranked in the *Miners* example. It also means that **DECLINE** will be top-ranked relative to **coarse**, but that **ACCEPT&WRITE** will be top-ranked relative to **fine**.²⁴

²⁴To be clear, we have only suggested that alternatives are ordered by expected value because this idea is fairly simple, yields intuitive verdicts for the range of cases we consider, and allows us to provide a concrete implementation of our general framework. Officially,

Just like Kratzer’s analysis, Account 1 makes Inheritance valid: if all of the best alternatives entail p , and p entails q , then all the best alternatives will entail q as well. So, no matter what the set of alternatives and ordering is, conjunctions such as (3a) (‘Prof. Procrastinate has to accept and write, but he doesn’t have to accept’) can never be true. This provides a straightforward explanation for their infelicity.

But Account 1 doesn’t yet explain the contrast between (1a) and (1b); why the former sounds false in the *Procrastinate* scenario, but the latter sounds true:

- (1) a. Prof. Procrastinate ought to accept.
- b. Prof. Procrastinate ought to accept and write.

We have claimed that there are two relevant sets of alternatives in the *Procrastinate* scenario: **coarse** and **fine**. Given a decision-theoretic ordering, (1a) is false relative to **coarse**, and (1b) is true relative to **fine**. However, the problem is that (1a) is *true* relative to **fine**, and (1b) is *false* relative to **coarse** (since ACCEPT doesn’t entail *Procrastinate accepts and writes*). So, what needs to be explained is why (1a) tends to be evaluated relative to **coarse**, but (1b) is evaluated relative to **fine**. We try to do this in the next two subsections.

3.2 Representation

Let us begin by considering why (1b) is evaluated relative to **fine** rather than **coarse**. The basic idea is that ‘ought’ claims carry certain presuppositions which force (1b) to be evaluated relative to the more fine-grained set of alternatives. We spell this out below.

Recall that **coarse** = {ACCEPT, DECLINE}. The complement in (1b) expresses the following proposition: *Procrastinate writes and accepts*. Intuitively, there is a mismatch here between this proposition and the distinctions made salient by **coarse**. More precisely, *Procrastinate writes and accepts* is

all we are committed to is that in the contexts of interest, the ordering over alternatives has a decision-theoretic aspect which allows us to distinguish between the alternatives. We want to leave it open whether the ordering is *always* determined by decision-theoretic considerations. For instance, the so-called “objective” sense of ‘ought’ might call for a ranking that departs from notions such as expected value. (See (Cariani, 2016b) for relevant discussion.) Moreover, we want to leave it open that in the cases of interest, the “good news value” of an alternative is better represented by using expected value in more sophisticated ways, or even by adopting a more sophisticated decision theory, e.g. prospect theory (Kahneman & Tversky, 1988) or risk-weighted utility theory (Buchak, 2013). Finally, we note that the background information state used to formulate the relevant decision-theoretic notion shouldn’t necessarily be taken to be the agent’s beliefs or credences. For example, see the game show cases in §3.3.

too specific, it cuts logical space finer than **coarse** allows. The thought is that this sort of disconnect between the prejacent of an ‘ought’ claim and the background set of alternatives should be ruled out. We can make this more precise through the following definition. Given a set of alternatives \mathcal{A} and proposition p , let us say that p is *represented by* \mathcal{A} just in case every alternative in \mathcal{A} either entails p or entails $\neg p$.²⁵ For instance, the proposition that Procrastinate accepts is represented by **coarse**, but the proposition that he accepts and writes is not, since ACCEPT neither entails this proposition nor its negation.

We propose adding a representation requirement to the semantics of ‘ought’ claims. One option is to add it as a regular truth condition. However, then a sentence such as (12) will be true relative to **coarse**:

(12) # Prof. Procrastinate doesn’t have to write the book review.

This is because DECLINE doesn’t entail *Procrastinate writes the review*. However, it is difficult to hear this sentence as acceptable.²⁶

Instead, we suggest capturing the representation requirement as a presupposition, or definedness condition on ‘ought’ claims. Our final entry then looks as follows:

Semantics for *ought*

‘ought p ’ is defined relative to $\langle w, \mathcal{A}, \succ \rangle$ only if p is represented by \mathcal{A}

If defined, ‘ought p ’ is true relative to $\langle w, \mathcal{A}, \succ \rangle$ iff for every $q \in \text{TOP}(\mathcal{A}, \succ)$: $q \subseteq p$

(12) is not predicted to be true on our account, since presuppositions project through negation (Chierchia & McConnell-Ginet, 2000). So, (12) is predicted to suffer from presupposition failure when evaluated relative to **coarse**.²⁷

²⁵Cf. Cariani’s (2013) notion of a proposition being “visible” with respect to a background partition of logical space. See §3.4 for further discussion of Cariani’s account.

²⁶Indeed, (12) seems bad even when it is uttered on the heels of a sentence that is plausibly evaluated relative to coarse, as in (13):

(13) # Prof. Procrastinate doesn’t have to accept, and he doesn’t have to write the book review.

²⁷For convenience, we assume that presupposition failure has a semantic effect. But our general approach is compatible with pragmatic accounts of presupposition on which presupposition failure affects assertability rather than semantic value (Schlenker, 2009). In this setting, there is every prospect for maintaining that Inheritance is classically valid (cf. the discussion immediately below).

Note that on this semantics, it is still the case that configurations of the form ‘ought p but not ought q ’, where p entails q , can never be true. This is because although our account doesn’t make Inheritance classically valid, it is still *Strawson valid* (von Fintel, 1999). Essentially, an argument from a set of sentences Γ to a sentence ψ is Strawson valid just in case whenever all of the $\varphi \in \Gamma$ and ψ are defined, if all of the $\varphi \in \Gamma$ are true, then ψ must be true as well.²⁸ So, sentences of the form ‘ought p but not ought q ’, where p entails q , will always either be false or undefined.²⁹

Now, it is generally accepted that hearers tend to interpret sentences so as to avoid presupposition failure. That is, within reasonable limits, hearers will make the assumptions necessary to avoid undefinedness.³⁰ For instance, even if I suspect that you have no siblings, I will come to assume that you have a sister once I’ve heard you utter (14):

(14) I have to fetch my sister from the airport tomorrow.

That is, the presupposition triggered by the possessive ‘my sister’, namely that I have a sister, is taken to hold when evaluating (14).

We want to suggest that something similar happens when we interpret (1b) (‘Prof. Procrastinate ought to accept and write’) in context. To repeat, there are two relevant sets of alternatives in the *Procrastinate* scenario, namely **coarse** and **fine**. We have seen that (1b) fails to be defined relative to **coarse**, since the proposition *Procrastinate accepts and writes* isn’t represented by this set of alternatives. But (1b) *is* defined relative to **fine**. Assuming that hearers will try to avoid interpretations that result in presupposition failure, they will then be moved to evaluate (1b) relative to **fine** rather than **coarse**. This is exactly the result that we wanted.

Note that the drive to avoid presupposition failure still doesn’t clarify why (1a) should be evaluated relative to **coarse**, since this sentence is defined

²⁸Strawson validity has been used to account for a range of natural language phenomena, e.g., see (von Fintel, 1999; Cariani & Goldstein, 2018; Mandelkern, 2020). That said, one can raise questions about how exactly the notion of Strawson validity bears on intuitions about natural language inferences. For instance, Strawson entailment fails to be transitive (Dorr & Hawthorne, 2018). We’re sympathetic to these concerns, but we don’t think that they should trouble us here. Most importantly, our use of Strawson validity is fairly constrained, since we’re essentially only interested in entailments from a single sentence schema.

²⁹To be clear, the infelicity of examples such as (3a) (‘Prof. Procrastinate has to accept and write, but he doesn’t have to accept’) doesn’t simply follow from the claim that they are Strawson contradictions on a uniform interpretation. Rather, their infelicity follows from that fact in combination with the claim that mid-sentence contexts shifts in these sentences are difficult to achieve (see the discussion of context shifting in §2.2).

³⁰This process is usually discussed under the heading “accommodation”. The basic phenomenon goes back at least to Karttunen (1974). See (von Fintel, 2004) for a more recent discussion.

relative to **fine** as well. Further arguments must be proffered in order to account for this asymmetry. We turn to these next.

3.3 Alternatives and immediate control

We want to start by getting clearer about the data surrounding examples such (1a), and what precisely needs to be explained. We have already seen that this sentence is unacceptable when uttered in the original *Procrastinate* scenario. But it is important to observe that the example improves when uttered in conjunction with (1b). For instance, consider the following:

- (15) a. Prof. Procrastinate ought to accept and write, so he ought to accept.
- b. Prof. Procrastinate ought to accept, since he ought to accept and write.

These seem basically fine to us, and certainly much better than (1a) uttered on its own. Also consider a variant of the *Procrastinate* scenario where your friend Pete is completely ignorant of the goodness/badness of Procrastinate's possible actions:

- (16) a. Pete: I know that Procrastinate can decline, accept and not write, or accept and write. What should he do?
- b. You: He ought to accept, because accepting and writing is best.

Again, we don't find it difficult to hear your reply in (16b) as true. What the above examples show is that (1a) isn't *always* unacceptable, and can be felicitous when uttered in the right setting. So, what needs to be explained is why (i) (1a) sounds bad when uttered in the most natural reconstructions of the *Procrastinate* scenario, but (ii) (1a) improves when uttered in configurations such as (15) and (16b).³¹

In §3.1 we proposed that the alternatives relevant for the evaluation of an 'ought' claim are tied to the actions that the relevant agent(s) can perform, their *decision problem*. We also mentioned that decision problems become more fine-grained over time, as the space of possible actions available to the agent increases. In order to account for (i), we suggest that the relevant alternatives tend to reflect what is in the agent's *immediate control*. That

³¹Note that the examples in (15) and (16b) pose a further problem for Inheritance-rejecting theories such as those considered in §2.1. For if (1a) is false in context, why should this sentence become more acceptable when uttered in conjunction with (1b)? It's unclear.

is, the alternatives used to evaluate a deontic claim are, *ceteris paribus*, derived from the agent's immediate courses of action, their short-term decision problem. If that's correct, then we should expect that (1a) will, by default, be interpreted relative to **coarse** rather than **fine**, since Procrastinate's immediate options are simply accepting and declining. And as we've seen, (1a) is *false* relative to **coarse**, which would explain its infelicity.

That the operative alternatives tend to align with what is in an agent's immediate control can be supported by reflecting on the general point and purpose of uttering deontically modalized claims. These claims are often uttered with an eye to advising agents of their best course of action, and aiding them in the the process of deliberation. For computational reasons, it is prudent for those who are providing advice to focus on the earlier branches of the agent's decision tree. The number of possible actions available to an agent grows large over time, and if presented with a complete course of action that spans the entire tree, it is easy for agents to become overwhelmed. Think about providing directions to a tourist: in order to maximize the chances of uptake, it is best to disclose the directional information in a step-by-step fashion; the tourist could easily become bewildered if presented with a complete route to their destination all in one go. Another point in this vicinity involves motivation: it is plausible that agents tend to be more motivated to finish complex tasks if they are broken down into smaller, discrete actions.³² If that's right, then it would make sense for deontic claims to be evaluated relative to a comparatively short-term temporal horizon.

There are also empirical reasons to think that alternatives tend to reflect what is in the agent's immediate control. Indeed, to our minds one of the strongest considerations in favor of this idea is that it makes the following prediction: it predicts that if we started to alter cases such as *Procrastinate* so that all of the actions available to the agent could be carried out at a single point in time, examples such as (1a) should start to sound more acceptable. It is difficult to test this in the actual *Procrastinate* case without the example becoming fairly contrived. But consider the following scenario:

Diachronic Treaty: You are a high government official for the country of Arun. Arun's diplomats have negotiated treaties with Arun's neighbors Bethesda and Carmill. Carmill is much stronger than Bethesda. Signing treaties with both neighbors would bring great wealth and prosperity to Arun, while signing with neither will have no effect. Signing a treaty with Bethesda but not Carmill will inevitably lead to war with Carmill, and the eventual ruin of Arun. Today you must choose whether to sign

³²This psychological tendency is reflected in the following well-known question and answer pair: How do you eat an elephant? One bite at a time!

a treaty with Bethesda. Tomorrow you must choose whether to sign a treaty with Carmill. You find Carmill's king unbearable, and detest facilitating the development of his kingdom. So, even if you sign with Bethesda today, there's a decent chance that you will not be able to bring yourself to sign with Carmill tomorrow.

(17) You ought to sign the treaty with Bethesda.

Our judgment is that (17) patterns with (1a): (17) is infelicitous in context. But now consider the following variant of the case:

Synchronic Treaty: As in *Diachronic Treaty*, but now the decision to sign with Bethesda and Carmill is made at the same time. Depending on which of four buttons you press, your electronic signature will be placed on both treaties, exactly one of the treaties, or none of them.

(18) You ought to sign the treaty with Bethesda.

To our ears, and those of our informants, (18) improves considerably here. Indeed, the sentence just sounds true. This is neatly accounted for by the idea that alternatives are "chunked" with respect to what is in an agent's immediate control. In *Diachronic Treaty*, your short-term decision problem involves whether or not to sign with Bethesda. Relative to this set of alternatives, (17) is false, since not signing with Bethesda has a higher expected value than signing. But in *Synchronic Treaty*, your short-term decision problem involves every configuration of accepting and declining treaty agreements. Relative to this set of alternatives, (18) is true, since signing with both Bethesda and Carmill has highest expected value.

Here is another example to a similar effect:

Diachronic Game Show: Alex is on a game show. He has to push two buttons consecutively, with a minute pause between pushes. He has a choice of the red button or the blue button for his first push, and a choice of the white button or the black button for his second push. If Alex pushes the red button first, he'll get \$1000 regardless which button he pushes after. If he pushes the blue button first, he'll either get \$1100 if he pushes the black button second, or \$0 if he pushes the white button second. Alex is entirely ignorant of the payoffs associated with the button pushes.

(19) Alex ought to push the blue button first.

Again, (19) patterns with (1a): (19) is infelicitous in context. But now consider the following variant of the case:

Synchronic Game Show: Alex is on a game show. He has to push two buttons at the same time: one with his left hand and one with his right hand. He has a choice of the red button or the blue button for his left hand push, and a choice of the white button or the black button for his right hand push. If Alex pushes the red button with his left hand, he'll get \$1000 regardless which button he pushes with his right hand. If he pushes the blue button with his left hand, he'll either get \$1100 if he pushes the black button with his right hand, or \$0 if he pushes the white button with his right hand. Alex is entirely ignorant of the payoffs associated with the button pushes.

(20) Alex ought to push the blue button with his left hand.

(20) sounds true here. Again, this can be explained if we assume that alternatives tend to reflect what is in an agent's immediate control. In *Diachronic Game Show*, Alex's short-term decision problem involves whether to push the red button or the blue button. Relative to this set of alternatives, (19) is false, since pushing the red button has a higher expected value than pushing the blue button. But in *Synchronic Game Show*, Alex's short-term decision problem involves pushing every button configuration. Relative to this set of alternatives, (20) is true, since pushing blue-black has highest expected value.

Further support for a connection between the operative alternatives and what is in an agent's control comes from the *Miners* case presented in §3.1:

Miners: Ten miners are trapped either in shaft A or in shaft B, but we do not know which. Flood waters threaten to flood the shafts. We have enough sandbags to block one shaft, but not both. If we block one shaft, all the water will go into the other shaft, killing any miners inside it. If we block neither shaft, both shafts will fill halfway with water, and just one miner, the lowest in the shaft, will be killed.

(11) We ought to block neither shaft.

As we observed, (11) is naturally heard as true here. This can be explained if we suppose that it is evaluated relative to the set $\mathcal{A}_1 = \{\text{NEITHER}, \text{A}, \text{B}\}$, since NEITHER has higher expected value than the other alternatives. But

consider the set of alternatives $\mathcal{A}_2 = \{\text{NEITHER}, \text{I}, \bar{\text{I}}\}$, where I is the proposition that we block the shaft the miners are in, and $\bar{\text{I}}$ is the proposition that we block shaft the miners aren't in. Clearly, I has highest expected value among these alternatives. So, if (11) was evaluated relative to \mathcal{A}_2 , then the sentence would be *false*. Why isn't (11) evaluated relative to something like \mathcal{A}_2 ? We suggest that this is because \mathcal{A}_2 doesn't reflect what's in our control. We can't choose whether or not to block the shaft which contains the miners, since we don't know which shaft they're in. To be clear, this is not to say that sets such as \mathcal{A}_2 are *never* relevant for evaluating deontic claims in *Miners*. For instance, it is easy to hear (21) as true:

(21) We ought to block the shaft that the miners are in.

This claim won't be defined relative to \mathcal{A}_1 . So, for reasons given in §3.2, (21) will tend to be evaluated relative to something like \mathcal{A}_2 .³³

As for (ii), we have been careful to frame the preference for short-term decision problems as a default, rather than a rigid rule. So, we predict that it could be overridden in the right circumstances, e.g. in order to avoid presupposition failure. Given our assumption that **coarse** and **fine** are the relevant sets of alternatives in the *Procrastinate* scenario, we then have a fairly straightforward explanation for why 'Procrastinate ought to accept' sounds fine in, e.g. (15a) ('Prof. Procrastinate ought to accept and write, so he ought to accept'). We have seen that the accept-and-write conjunct is undefined relative to **coarse**. So, in order to ensure that this conjunct is defined, hearers will evaluate these conjunctions relative to **fine**. But (1a) is *true* relative to **fine**.³⁴

³³Some might worry that our explanation for why (1a) is unacceptable contravenes a compelling principle involving speaker interpretation: when possible, hearers will interpret utterances to be true. Given this principle, our explanation might seem surprising, since hearers are able to evaluate (1a) relative to a set of alternatives on which it comes out true, namely **fine**. However, care must be taken when applying interpretation principles such as these. The psycholinguistics literature is replete with examples where on the face of it hearers should easily be able to access true interpretations, and yet fail to do so. For instance, consider pronoun resolution (Kehler & Rohde, 2013):

(22) Norm lent his car to his brother's girlfriend. He doesn't own one.

Subjects tend to judge (22) to be false, and question how it is that Norm could both lend a car and fail to own one. But one might antecedently have expected that 'he' could easily be resolved to Norm's brother's girlfriend, which could have led to a true interpretation.

³⁴A reviewer wonders why presupposition accommodation does not induce a mid-sentence context shift in (3b) ('Prof. Procrastinate doesn't have to accept, but he has to accept and write') so that the first conjunct is evaluated relative to **coarse** and the second conjunct is evaluated relative to **fine**, in which case the conjunction would be true. The reviewer motivates this possibility by considering examples such as (23) (styled after (Lewis, 1979)):

Note that our account of why (1a) is unacceptable, but (1b) is acceptable, is similar to the explanations given by von Stechow and Bronckmann & Dowell discussed in §2.2, since it also involves appealing to a context shift. However, there are important differences. For one thing, we suggested that the most plausible way of developing von Stechow's response predicts that (1a) should be acceptable in the following variant of the *Procrastinate* scenario (repeated from above):

Procrastinate 2: Prof. Procrastinate is invited to review a book on which he is the only fully qualified specialist on the planet. Procrastinate is generally reliable, but sometimes fails to bring projects to completion. In particular, if Procrastinate accepts to review the book, it is quite likely (but not certain) that he will end up writing the review. If Procrastinate accepts but doesn't end up writing the review, the consequences will be catastrophic: his actions will precipitate a crisis in tertiary education that will culminate in several universities closing and thousands of academics losing their livelihoods. If Procrastinate declines, someone less qualified than him will end up writing the review, but there will be no crisis.

But (1a) still sounds false here, since the costs of Procrastinate accepting and failing to write the review are enormous. This is what we predict, since

(23) The dog bit the other dog.

The idea is that the presupposition carried by the definite 'the other dog' triggers domain expansion, ultimately yielding a coherent interpretation. However, we are doubtful that the pressure to accommodate presuppositions can alone induce intersentential context shifts. For instance, in (23) the adjective 'other' is crucial for inducing the shift here. For observe that it is much harder to access a good reading of (24):

(24) ?? The dog bit the dog.

In fact, it is striking that even with contextual clues, e.g. if we successively pointed at each dog, (24) still sounds bad. A further example in this vein involves possessives:

(25) ?? Bill's dog bit Bill's dog.

Even if we thought that Bill had at least two dogs, (25) would still be infelicitous. Or consider (26):

(26) # There are no dogs but the dog is happy.

Clearly, the pressure to accommodate the presupposition triggered by the definite isn't sufficient to achieve domain expansion. Similar examples can be constructed using different presupposition triggers, e.g. 'both dogs'. In short, expressions such as 'other' play an essential role in inducing the context shift in (23). But there are no expressions similar to 'other' that feature in (3b) which could facilitate a context shift through presupposition accommodation.

Procrastinate’s short-term decision problem still just involves accepting or declining. And given the high costs of him accepting and not writing, declining has high expected value.

As for Bronfman & Dowell, they predict that the examples in (10) should have true readings:³⁵

- (10) a. Prof. Procrastinate ought to/has to not accept and write.
b. Prof. Procrastinate ought to/has to do something other than accept and write.

By contrast, we predict that these examples will fail to be true on either **coarse** or **fine**. For instance, relative to **coarse**, (10a) will be undefined, since the proposition *Procrastinate doesn’t accept and write* isn’t represented by this set of alternatives. And relative to **fine**, the sentence will be false, since accepting and writing is the top-ranked alternative.

To sum up, we have proposed an account of ‘ought’ claims which patterns broadly with Kratzer’s analysis: ‘ought’ operates as a quantifier over a restricted domain of possibilities. However, on our account the objects in this domain are propositions, or alternatives, rather than worlds. Moreover, ‘ought’ triggers presuppositions which impose constraints on the relationship between the prejacent and the background set of alternatives. In combination with auxiliary principles governing the choice of alternatives in context, we argued that our proposal resolves the tension involving Inheritance that we set out to explain.³⁶

³⁵Bronfman & Dowell also predict that ‘We ought not block the shaft the miners are in’ has a true interpretation in the *Miners* scenario. But as with the examples in (10), it is very challenging to hear this sentence as good.

³⁶A reviewer discusses the following case, where you are deciding what to do with some leftovers. You can eat the leftovers now, or you can eat them tomorrow. There is also a very small chance that you forget about the food and let it rot in the fridge. Suppose that you aren’t very hungry now, so it would be a little bit better to eat the leftovers tomorrow. The reviewer maintains that one can set up the utilities and probabilities in this case so that our account predicts that (27) should be true relative to $\{\text{NOW}, \overline{\text{NOW}}\}$, where NOW is the proposition that you eat the food now, and $\overline{\text{NOW}}$ is the proposition that you don’t eat the food now.

- (27) You ought to eat the leftovers now.

The reviewer argues that this is the wrong prediction, since (27) is unequivocally false in this context.

In response, we think that the reviewer’s case is rather underdescribed as stated. And when one fills in the details, we think that our account makes the correct predictions. Let us suppose that the probability of you forgetting about the leftovers conditional on you not eating them now is $\frac{1}{1000}$. It is helpful to assign utilities to the outcomes in line with the reviewer’s description of the case: let us say that $\text{EV}(\textit{eat now}) = 10$, and $\text{EV}(\textit{wait and eat tomorrow}) = 12$ (we make the difference in utilities small because this is what the reviewer

We end this section by comparing our account to the theories of Cariani (2013, 2016b).

3.4 Cariani (2013, 2016b)

We aren't the first to marry a Kratzer-style entry with alternative-sensitivity. In particular, Cariani develops two distinct proposals that are broadly in this vein. We briefly consider these approaches here. In short, neither (Cariani, 2013) nor (Cariani, 2016b) provides an adequate response to our target observations, but one could develop (Cariani, 2016b) so that the resulting account is similar to our own.

Cariani (2013) also maintains that 'ought' quantifies over a domain of alternatives, and he also imposes a representation requirement. However, his account differs from our own in two respects: (i) the representation condition is captured as a regular truth-condition, rather than a presupposition; and (ii) Cariani posits a "benchmark element" in the contextually determined ordering \succ , and requires that \lceil ought $p \rceil$ is true only if every p -entailing alternative is at least as good as this benchmark. Cariani's account is roughly the following:

stipulated). Then in order for $EV(\textit{eat now}) > EV(\textit{don't eat now})$, $EV(\textit{let leftovers rot}) < 1000 \cdot (10 - 12 \cdot \frac{999}{1000}) = -1988$. In other words, the outcome of letting the leftovers rot in the fridge has to be *very* bad in order for $EV(\textit{eat now}) > EV(\textit{don't eat now})$ given our background assumptions. But it is plausible that in most normal circumstances, letting the food rot *wouldn't* be so bad. For instance, you will still have enough money to buy other food, so you won't starve. Thus, in most normal circumstances, $EV(\textit{let leftovers rot}) \geq -1988$. In that case, our account predicts that (27) should be false.

Now suppose that the circumstances aren't normal. In what sorts of scenarios would $EV(\textit{let leftovers rot})$ be so low, i.e. smaller than -1988? One possible scenario is the following: suppose that the leftovers will produce a noxious, undetectable gas after two days of being left in the fridge, and the gas will cause all the other food in the fridge to go bad, causing hundreds of dollars of spoilage. In this context, our account predicts that (27) should be true. To our minds, this is the correct prediction. Even though you aren't particularly hungry, it is reasonable to eat the leftovers now rather than risk, e.g. the entire fridge being spoiled by not eating them now. Indeed, this choice is reflected in judgments involving relative betterness: a sentence such as 'It is better to eat the leftovers now rather than risk an entire fridge full of spoiled food by not eating them now' seems true in this scenario. Another way of putting the point is that this variant of the case essentially becomes a version of the Procrastinate scenario: even if it is likely that Procrastinate remembers to write the review, and even if accepting and writing is only a little bit better than declining, it is easy to hear 'Procrastinate ought not accept' as true if Procrastinate accepting but not writing would be catastrophic, e.g. if it would precipitate a crisis in tertiary education that will culminate in several universities closing and thousands of academics losing their livelihoods. (The *Procrastinate 2* scenario from §2.2 also illustrates that the event of Procrastinate writing being unlikely is inessential to getting the relevant pattern of judgments going.) In short, when the reviewer's case is made more specific, and utilities are clearly assigned, our account appears to make the correct predictions.

Cariani’s (2013) semantics for *ought*

$\lceil \text{ought } p \rceil$ is true relative to $\langle w, \mathcal{A}, \succ \rangle$ iff

- (i) for every $q \in \text{TOP}(\mathcal{A}, \succ)$: $q \subseteq p$; and
- (ii) for every $q \in \mathcal{A}$ such that $q \subseteq p$, $q \succeq$ the benchmark; and
- (iii) p is represented by \mathcal{A}

Granted plausible assumptions, this account predicts that (1a) (‘Prof. Procrastinate ought to accept’) should be false, and (1b) (‘Prof. Procrastinate’) should be true. For instance, suppose that the relevant set of alternatives is **fine**, and the ordering looks as follows:

$$\text{ACCEPT\&WRITE} \succ \text{DECLINE} \succ \text{benchmark} \succ \text{ACCEPT\&}\overline{\text{WRITE}}$$

Then clearly (1b) will be true, but (1a) will be false since $\text{ACCEPT\&}\overline{\text{WRITE}}$ entails that Procrastinate accepts, but it is below the benchmark.

However, the truth-conditions this account generates are much too strong. For instance, suppose that you are choosing between beef, chicken, prawns, and tofu. You like the beef and the tofu, but we come to learn that both the chicken and the prawns are poisoned.

(28) You ought not order the chicken, and you ought not order the prawns.

(28) is impeccable here. But Cariani’s account predicts that it is false on the most natural set of alternatives, i.e. $\mathcal{A} = \{\text{BEEF}, \text{CHICKEN}, \text{PRAWNS}, \text{TOFU}\}$. Presumably, the ordering over these alternatives looks as follows, with both CHICKEN and PRAWNS below the benchmark:

$$\text{BEEF/TOFU} \succ \text{benchmark} \succ \text{CHICKEN/PRAWNS}$$

Then the first conjunct in (28) will be false, since PRAWNS entails $\neg \text{You order chicken}$, but it is below the benchmark. One might try to maintain that the first conjunct of (28) is evaluated relative to a more coarse-grained set of alternatives on which it comes out true, e.g. $\{\text{CHICKEN}, \overline{\text{CHICKEN}}\}$. But then assuming that mid-sentence context shifts are difficult to achieve, the second conjunct in (28) will be false.^{37,38}

³⁷This response also predicts that an utterance of (28) should degrade when each of the four meal alternatives are made salient prior to the utterance. But this doesn’t seem to be the case.

³⁸Examples such as (28) also pose a problem for the actualist account presented in §2.1. Suppose that the chicken and the prawns are equally poisonous. Then ‘You ought not order the chicken’ will be false if you actually end up ordering the prawns. This is because the closest world where you don’t order chicken will be the actual world, which is a world where you order the prawns. And the actual world will be just as bad as the closest world where you order chicken.

Moreover, this account leaves our second observation unexplained. It makes Inheritance straightforwardly invalid, and predicts that conjunctions such as (3a) should be true relative to **fine**:

- (3a) # Prof. Procrastinate has to accept and write, but he doesn't have to accept.

In this respect, the account does no better than the other Inheritance-rejecting views considered in §2.1.³⁹

As for the account of Cariani (2016b), this proposal is closer to Kratzer's original semantics in that the alternatives are essentially only used to select the worlds in the domain of quantification. The entry looks as follows:⁴⁰

Cariani's (2016b) semantics for *ought*

⌈ought p ⌋ is true relative to $\langle w, \mathcal{A}, \succ \rangle$ iff every $w \in \text{TOP}(\mathcal{A}, \succ)$ is a p -world

This entry is essentially equivalent to our initial entry presented in §3.1.⁴¹ It makes Inheritance valid, and so explains why conjunctions such as (3a) are abominable: they are always false. However, it doesn't account for our first cluster of observations, and is missing some crucial ingredients. Even if we assume that (1a) is false relative to **coarse**, and (1b) is true relative to **fine**, (1a) will then be *true* relative to **fine**, and (1b) will be *false* relative to **coarse**. So, what needs to be clarified is why (1a) is evaluated with respect to **coarse**, and (1b) is evaluated relative to **fine**. A major goal of this paper has been to provide an explanation. So, although we think that (Cariani, 2016b) is on the right track, it leaves some of our central phenomena unaccounted for.⁴²

Overall, neither of Cariani's entries—on their own—provides us with a satisfying explanation of our target observations.

³⁹See (Bronfman & Dowell, 2018) for several other criticisms of Cariani's (2013) semantics.

⁴⁰Cariani *et al.* (2013) present a similar entry. It is worth mentioning that Cariani (2016b) is primarily interested in trying solve a different set of problems from those that arise in the *Procrastinate* scenario, namely ones surrounding the way that deontic claims interact with probabilities and conditional operators. It is beyond the scope of this paper to consider how our theory could be used to address these latter puzzles.

⁴¹Strictly speaking, Cariani orders the alternatives in \mathcal{A} by considering their intersection with a background set of worlds, or *information state*. We could have used an information state parameter in a similar way on our account, but since nothing hangs on it for our purposes, we have opted for a simpler representation.

⁴²Note that there is independent reason for Cariani to impose a representation requirement on deontic modals. He treats modals such as 'may' and 'permit' as duals of 'ought'/'has to':

Cariani's (2016b) semantics for *may*

⌈may p ⌋ is true relative to $\langle w, \mathcal{A}, \succ \rangle$ iff some $w \in \text{TOP}(\mathcal{A}, \succ)$ is a p -world

4 Conclusion

We’ll close with a few remarks about the general status of Inheritance. Theorists have challenged Inheritance on a variety of grounds, not just for reasons stemming from scenarios such as *Procrastinate*.⁴³ In this respect, our aims in this paper have been relatively modest. We haven’t attempted to provide a complete defense of Inheritance here. Instead, we focused on what we take to be one of the most puzzling phenomena involving Inheritance, and tried to show that a satisfying resolution of this puzzle can be provided. How our account can be deployed to respond to other concerns involving Inheritance requires further investigation. But we are hopeful that our general approach is on the right track, and will prove useful to theorists interested in the semantics and pragmatics of deontic modals.

References

- Bronfman, Aaron, & Dowell, J. L. 2018. The Language of “Ought”, and Reasons. *In*: Star, Daniel (ed), *The Oxford Handbook of Reasons and Normativity*. Oxford University Press.
- Buchak, Lara. 2013. *Risk and Rationality*. Oxford University Press.
- Cariani, Fabrizio. 2013. ‘Ought’ and Resolution Semantics. *Noûs*, **47**(3), 534–558.

But then examples such as those in (29) are predicted to be true in *Procrastinate*:

- (29) a. Prof. Procrastinate may accept and write, and murder someone.
b. Prof. Procrastinate may decline and murder someone.

(29a) will be true relative to **fine**, since there is some world in **ACCEPT&WRITE** where Procrastinate accepts, writes, and murders someone. Similarly, (29b) will be true relative to **coarse**, since there is some world in **DECLINE** where Procrastinate declines and murders someone. But it is very hard to access a true reading of either example. If deontic claims carried a representation presupposition, then neither sentence would be defined.

⁴³One challenge comes from Ross’s Puzzle, which involves the observation that (30a) doesn’t seem to entail (30b):

- (30) a. You ought to mail the letter. $\not\Rightarrow$
b. You ought to mail the letter or burn it.

It is worth observing that although the inference from (30a) to (30b) is Strawson-valid on our approach, (30b) will fail to be defined on the most natural set of alternatives, i.e. $\{\text{MAIL}, \overline{\text{MAIL}}\}$. This could go some way to accounting for the intuitions that undergird Ross’s Puzzle. Crnič (2011) and von Fintel (2012) also suggest that the puzzle involves certain effects that arise from the special properties of disjunction. How exactly these properties should be captured is a matter of active investigation—see the discussion in (von Fintel, 2012) for some choice points. Also see (Cariani, 2021) for a good overview of the issues surrounding Inheritance.

- Cariani, Fabrizio. 2016a. Consequence and Contrast in Deontic Semantics. *Journal of Philosophy*, **113**(8), 396–416.
- Cariani, Fabrizio. 2016b. Deontic Modals and Probability: One Theory to Rule Them All? *In*: Charlow, Nate, & Chrisman, Matthew (eds), *Deontic Modality*. Oxford University Press.
- Cariani, Fabrizio. 2021. Deontic Logic and Natural Language. *In*: Gabbay, Dov, van der Meyden, Ron, Horty, John, Parent, Xavier, & van der Torre, Leandert (eds), *The Handbook of Deontic Logic (Vol. II)*. College Publications.
- Cariani, Fabrizio, & Goldstein, Simon. 2018. Conditional Heresies. *Philosophy and Phenomenological Research*.
- Cariani, Fabrizio, Kaufmann, Magdalena, & Kaufmann, Stefan. 2013. Deliberative Modality Under Epistemic Uncertainty. *Linguistics and Philosophy*, **36**(3), 225–259.
- Chierchia, Gennaro, & McConnell-Ginet, Sally. 2000. *Meaning and Grammar: An Introduction to Semantics*. MIT Press.
- Condoravdi, Cleo. 2002. Temporal Interpretation of Modals: Modals for the Present and for the Past. *In*: David Beaver, Stefan Kaufmann, Brady Clark, & Casillas, Luis (eds), *The Construction of Meaning*.
- Crnić, Luka. 2011. *Getting even*. Ph.D. thesis, MIT.
- Dorr, Cian, & Hawthorne, John. 2018. *If...: A theory of conditionals*. NYU and USC.
- Finlay, Stephen. 2019. *Ends, Probabilities, and Deontic Paradoxes*. Handout for Kioloa.
- von Fintel, Kai. 1999. NPI Licensing, Strawson Entailment, and Context Dependency. *Journal of Semantics*, **16**(2), 97–148.
- von Fintel, Kai. 2004. Would You Believe It? The King of France is Back! (Presuppositions and Truth-Value Intuitions). *In*: Reimer, Marga, & Bezuidenhout, Anne (eds), *Descriptions and Beyond*. Clarendon Press.
- von Fintel, Kai. 2012. The best we can (expect to) get? challenges to the classic semantics for deontic modals. Central APA, Chicago, IL.
- von Fintel, Kai, & Iatridou, Sabine. 2008. *How to Say Ought in Foreign: The Composition of Weak Necessity Modals*. Dordrecht: Springer Netherlands. Pages 115–141.
- Frank, Annette. 1997. *Context Dependence in Modal Constructions*. Ph.D. thesis, Universität Stuttgart.
- Goble, Lou. 1996. Utilitarian Deontic Logic. *Philosophical Studies*, **82**(3), 317–357.
- Iatridou, Sabine, & Zeijlstra, Hedde. 2013. Negation, Polarity, and Deontic Modals. *Linguistic Inquiry*, **44**(4), 529–568.
- Jackson, Frank. 1985. On the Semantics and Logic of Obligation. *Mind*, **94**(374), 177–195.
- Jackson, Frank. 1991. Decision-Theoretic Consequentialism and the Nearest and Dearest Objection. *Ethics*, **101**(3), 461–482.

- Jackson, Frank. 1997. Which Effects. *Pages 42–53 of: Dancy, J. (ed), Reading Parfit*. Blackwell.
- Jackson, Frank, & Pargetter, Robert. 1986. Oughts, Options, and Actualism. *Philosophical Review*, **95**(2), 233–255.
- Jeffrey, Richard C. 1965. *The Logic of Decision*. University of Chicago Press.
- Kahneman, Daniel, & Tversky, Amos. 1988. *Prospect theory: An analysis of decision under risk*. New York, NY, US: Cambridge University Press.
- Karttunen, Lauri. 1974. Presupposition and Linguistic Context. *Theoretical Linguistics*, **34**(1), 181–.
- Kehler, Andrew, & Rohde, Hannah. 2013. Aspects of a theory of pronoun interpretation. *Theoretical Linguistics*, **39**(3-4), 295 – 309.
- Kennedy, Christopher. 2007. Vagueness and Grammar: The Semantics of Relative and Absolute Gradable Adjectives. *Linguistics and Philosophy*, **30**(1), 1–45.
- Kolodny, Niko, & 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.
- Kratzer, Angelika. 1981. The Notional Category of Modality. In: Eikmeyer, H. J., & Rieser, H. (eds), *Words, worlds, and contexts: New approaches in world semantics*. Berlin: de Gruyter.
- Kratzer, Angelika. 1991. Modality. In: von Stechow, Arnim, & Wunderlich, D (eds), *Semantik/semantics: An international handbook of contemporary research*. Berlin: de Gruyter.
- Kratzer, Angelika. 2012. *Modals and Conditionals*. Oxford University Press.
- Lassiter, Daniel. 2017. *Graded Modality: Qualitative and Quantitative Perspectives*. Oxford: Oxford University Press.
- Lewis, David. 1979. Scorekeeping in a Language Game. *Journal of Philosophical Logic*, **8**(1), 339–359.
- Lewis, Karen S. 2016. Elusive Counterfactuals. *Noûs*, **50**(2), 286–313.
- Mandelkern, Matthew. 2020. *If p, then p!* Oxford University.
- Moss, Sarah. 2012. On the Pragmatics of Counterfactuals. *Noûs*, **46**(3), 561–586.
- Parfit, Derek. 1988. *What We Together Do*.
- Prichard, H. A. 1932. Duty and Ignorance of Fact. *Philosophy*, **8**(30), 226–228.
- Regan, Donald H. 1980. *Utilitarianism and Co-Operation*. Oxford University Press.
- Ross, W. D. 1939. *Foundations of Ethics: The Gifford Lectures Delivered in the University of Aberdeen, 1935–*. Oxford University Press.
- Schlenker, Philippe. 2009. Local Contexts. *Semantics and Pragmatics*, **2**(3), 1–78.
- Schroeder, M. 2011. Ought, Agents, and Actions. *Philosophical Review*, **120**(1), 1–41.
- Soames, Scott. 1986. Incomplete Definite Descriptions. *Notre Dame Journal of Formal Logic*, **27**(3), 349–375.
- Stanley, Jason, & Szabó, Zoltan Gendler. 2000. On Quantifier Domain Restriction. *Mind and Language*, **15**(2-3), 219–261.
- Zvolensky, Zsófia. 2002. Is a Possible-Worlds Semantics of Modality Possible? A Problem for Kratzer’s Semantics. In: Jackson, B (ed), *Proceedings of SALT 12*. Ithaca: Cornell.