

Deontic Reasoning Across Contexts^{*}

Justin Snedegar

University of St Andrews, St Andrews, Scotland
js280@st-andrews.ac.uk

Abstract. Contrastivism about ‘ought’ holds that ‘ought’ claims are relativized, at least implicitly, to sets of mutually exclusive but not necessarily jointly exhaustive alternatives. This kind of theory can solve puzzles that face other linguistic theories of ‘ought’, via the rejection or severe restriction of principles that let us make inferences between ‘ought’ claims. By rejecting or restricting these principles, however, the contrastivist takes on a burden of recapturing acceptable inferences that these principles let us make. This paper investigates the extent to which a contrastivist can do this.

Keywords: deontic modals · contrastivism · context-sensitivity · reasoning

This paper is about deontic reasoning, or reasoning with *oughts*. This is very plausibly the kind of reasoning deontic logicians aim to formalize. But I am not here directly concerned with formal logical systems. Rather, I am interested in natural language inferences we can make using the English word ‘ought’. In particular, I am concerned with whether a particular kind of theory of the meaning of ‘ought’ can explain some intuitive deontic inferences, which I take to be data that such a theory needs to explain.

Insights from deontic logic are relevant for developing and evaluating accounts of natural language deontic reasoning, and vice versa. Thus, a next step in this project would be to explore proposals in deontic logic, to see what they can teach us about analogous issues in natural language deontic reasoning. Unfortunately, though, all I can do here is flag places where this strategy, of drawing on lessons from deontic logic, seems promising. Near the end of the paper, I will suggest that recent work in preference-based deontic logic, in particular, offers an interesting avenue for future development of the picture I develop.

1 Cross-Context Deontic Reasoning

A simple, standard semantics for ‘ought’ holds that what you ought to do is what you do in all the best worlds. This simple picture leads to well-known puzzles of deontic reasoning, which have led to interesting complications of the

^{*} This version is forthcoming in the proceedings of DEON 2014. This project is ongoing.

semantics (and of the linguistic theory more generally) of ‘ought’. For example, the orthodox Kratzerian semantics for modals, developed in [1], is motivated in part by its ability to solve some of these puzzles. To simplify somewhat, Kratzer’s theory relativizes ‘ought’ (and other modal) claims to contextual parameters, including bodies of information and standards, and holds that the puzzles arise from ignoring shifts in these parameters. So what appear to be paradoxes really involve some kind of equivocation.

A different kind of complication is distinctive of what we can call *contrastivism* about ‘ought’—the thesis that ‘ought’ claims are always at least implicitly relativized to sets of alternatives. So to say that you ought to A is really, according to contrastivism, to say that you ought to A out of some set of alternatives.¹ The contrastivist holds that various puzzles about ‘ought’ arise from ignoring the contrast-sensitivity of ‘ought’ claims; so again, these puzzles, according to the contrastivist, are due to a kind of equivocation.

Contrastivism, and contextualist theories more generally, provide nice solutions to a variety of puzzles involving ‘ought’. But it is easy to overgenerate fallacious equivocations. Some deontic reasoning takes place across contexts. As the context shifts, the content of the relevant ‘ought’ claims will also shift, according to these contextualist theories. So the concern is that even deontic reasoning which seems unobjectionable will turn out to be fallacious. This is a general statement of the problem I take up in this paper.

Another way to bring out this point, and also a way to preview discussion later in the paper, is to focus on a particular inference rule that simple, non-contextualist theories of ‘ought’ validate, but that contextualist theories, including contrastivism, seem to invalidate.²

Inheritance: If p entails q , then if it ought to be that p , it ought to be that q .

This kind of inference is involved in many of the puzzles of deontic reasoning.

The trouble is that **Inheritance** also explains lots of good instances of deontic reasoning. Consider the following simple example:

- (1) I ought to buy milk.
- (2) Buying milk entails buying a dairy product.
- (3) So I ought to buy a dairy product.

(1)-(3) looks like an example of good deontic reasoning. If we accept **Inheritance**, we have an easy explanation for why it is good: it is just an instance of **Inheritance** (plus *modus ponens*). Once we reject this principle, though, we seem to lack an explanation.

The obvious move for a contextualist is to point out that she does not *reject* **Inheritance** but only *restricts* its application. Specifically, she restricts its application to deontic reasoning that takes place in a single context. So as long as

¹ For different versions of contrastivism, see [2–5].

² An analogous axiom in deontic logic is often called ‘inheritance’ or ‘necessitation’: If $\vdash p \rightarrow q$, then $\vdash Op \rightarrow Oq$. The principle I call ‘Inheritance’ below, though, is rather a schema for explaining natural language inferences using ‘ought’.

we remain in the same context—as long as the relevant contextual parameters don't shift—throughout the reasoning process encoded in (1)–(3), we can apply this restricted version of **Inheritance** to explain why this is good reasoning. The hope is then that there are not single-context versions of the puzzles.³

The trouble is that lots of deontic reasoning takes place across contexts. For example, if we follow Kratzer and hold that 'ought' claims are relativized to, among other things, bodies of information, we can imagine a situation in which premise (1) is accepted relative to one body of information, then some information is added (including, if the reasoner is sufficiently ignorant about milk, premise (2)), and finally the conclusion is accepted relative to this updated body of information. Cases like this can still be instances of good deontic reasoning, but they are not explained by the restricted version of **Inheritance**.⁴

Though I think that recapturing good cross-context inferences is a general issue that any contextualist theory should address, my focus here will be on contrastivism. Once we focus on contrastivism, the problem of recapturing good cross-context reasoning becomes the problem of explaining inferences between 'ought' claims that are relativized to different sets of alternatives. After introducing contrastivism in more detail in section 2, I will show, in section 3, that it is relatively easy to secure some of these inferences. In the remainder of the paper, though, I return to **Inheritance**-like inferences. More work must be done to secure these. The only attempt of which I am aware, from Cariani in [5], faces some problems. I will show that a contrastivist picture that is actually simpler in some ways than Cariani's picture can avoid these problems while still securing some attractive inferences.

2 Contrastivism

Contrastivism about 'ought' holds that 'ought' claims are relativized, at least implicitly, to sets of alternatives. As the alternatives shift, the truth of the 'ought' claims can shift, as well. According to the contrastivist, there will be some actions such that, whether you ought to perform them can vary with the alternatives. This is to say more than that the *availability* of other alternatives can affect whether you ought to perform a given action—nearly everyone would accept that. The distinctive contrastivist claim is that the particular alternatives to which we are comparing the action can affect whether or not you ought to perform the action. Here is a simple illustration:

(4) You ought to take the bus rather than drive your SUV.

³ In deontic logic, lots of work has been done to give restricted versions of the analogous principle, sometimes called 'necessitation', in order to avoid paradoxes of deontic logic. This literature is very large; [6] surveys attempts with an eye specifically to allowing for deontic conflicts. See also [7–9]. Later in the paper I will point out some similarities between contrastivist proposals and some proposals from this formal work.

⁴ See [10] for a similar complaint against contextualist formulations of an epistemic closure principle.

- (5) But it's not the case that you ought to take the bus rather than ride your bike.

Putting things explicitly in terms of sets of alternatives, instead of 'rather than' claims:

- (4') You ought to take the bus out of {take the bus, drive your SUV}.
 (5') It's not the case that you ought to take the bus out of {take the bus, ride your bike}.

The contrastivist holds that both of these claims may be true, even if all three actions are available.

The alternatives may be provided explicitly, as in the 'rather than' ascriptions (4) and (5) above. More often, they will be provided implicitly, by the *question under discussion*—in particular by the *deliberative* question under discussion. Questions are sets of alternatives, with each alternative corresponding to a potential answer to the question. A *deliberative* question is a special kind of question—a question of what to do—with each alternative corresponding to some action that the relevant agent can perform.⁵

Contrastivism can solve puzzles about 'ought'. Jackson, in [3], uses the following example to motivate his contrastive view of 'ought'. Imagine the following dialogue taking place:

- A:** It ought to be that Lucretia used less painful poisons.
B: Oh no, it ought to be that she used painless poisons.
C: Oh no, it ought to be that she used political means rather than poison to obtain her ends.
D: Oh no, it ought to be that she never existed at all.
E: Oh no, it ought to be that she existed but made people happy.
F: Oh no, it ought to be that everyone was *already* happy.
 And so on.

As Jackson points out, each of these claims seems perfectly appropriate, when uttered. But according to the principle **Inheritance**, they are inconsistent. To focus on just one, suppose that B's claim is true. Then by **Inheritance**, it ought to be that Lucretia existed, since it ought to be that she used painless poisons, and her using painless poisons entails that she existed. But this contradicts D's claim.

Jackson's solution is to hold that "our everyday judgments of what ought to be are all relative to sets of (mutually exclusive but not necessarily jointly exhaustive) alternatives" (p. 180). As we move through the dialogue about Lucretia, the alternatives shift. A's claim is relative to a set like {she uses more painful poisons, she uses less painful poisons}, B's is relative to {she uses (less)

⁵ See [11, 12] for classic discussions of this conception of questions. See [13] and the references there for discussion of questions under discussion and their role in theories of communication. See [5] for a contrastivist theory of 'ought' developed in terms of deliberative questions.

painful poisons, she uses painless poisons}, C's is relative to {she uses (painless) poisons, she uses political means}, and so on. Since these claims are relativized to different sets of alternatives, there is in fact no inconsistency. Of course, if we were to evaluate an earlier claim relative to a later set of alternatives, we would get inconsistencies. But the data is that the claims seem perfectly appropriate *when made*. Once we consider, say, C's claim, B's claim no longer seems true. The contrastivist explanation is that once we consider C, we have in mind a set of alternatives relative to which B's claim is actually false.

This same kind of move provides a solution to the famous Good Samaritan paradox. This can be seen clearly by focusing on claim E above. 'It ought to be that she existed and made everyone happy' seems true, but she could only make everyone happy if some people were unhappy to begin with (let's suppose). So by **Inheritance**, the truth of E's claim would seem to imply the truth of 'It ought to be that some people were unhappy to begin with', which seems false. But again, by relativizing to sets of alternatives, we can block this inference. E's claim is only true relative to a set like {she existed and made people miserable, she didn't exist and some people were miserable, she existed and made everyone happy}. Relative to this set, it is not true that it ought to be that some people were unhappy to begin with. So the Good Samaritan inference doesn't seem to go through.

According to Jackson's contrastivist theory, it ought to be that p out of Q just in case p is the best out of Q . Or, to put things in agential terms (though Jackson does emphasize that he means to be talking about the 'ought to be' rather than the 'ought to do'—I will try to gloss over this issue here), you ought to A out of Q just in case A is the best alternative in Q .⁶ Other contrastivist theories offer more or less significant variations from this simple picture. So the contrastivist reconciliation of the claims above can be put like this: an action which is the best out of one set of alternatives need not be best out of a different set of alternatives.

3 Simple Inferences

One way to think of the contrastivist solutions is as blocking problematic inferences between 'ought' claims. For example, we cannot infer from A's claim one that would be inconsistent with B's, since this would involve a shift in the set of alternatives, and so an equivocation. This description of the solution fits with the broader contextualist move I discussed at the beginning of the paper. But it also highlights the problem of accounting for good cross-context deontic reasoning. It seems that we *can* make some inferences between 'ought' claims which are (according to the contrastivist) relativized to different sets of alternatives.

For example, suppose I decide that I ought to go to church out of {go to church, go to the bar, go to the office}. If going to the bar becomes irrelevant for some reason—maybe the person I wanted to meet there isn't going, or maybe

⁶ See [14] for a contrastivist theory that attempts to maintain a unified theory of the 'ought to do' and the 'ought to be'.

the bar doesn't open on Sundays—I should be able to infer straightforwardly that I ought to go to church out of the new set of alternatives, {go to church, go to the office}. The removal of going to the bar as a relevant alternative, when it wasn't what I ought to do anyway, shouldn't make a difference.

For another example, if I decide that we ought to go to church out of {go to church, go to the bar, go to the office} and you decide that we ought to go to church out of {go to church, go to the game, stay home}, then when we get together to talk about what to do, we should be able to straightforwardly infer that we ought to go to church out of {go to church, go to the bar, go to the office, go to the game, stay home}. If we come to the same conclusion about what we ought to do, even if we're considering different alternatives, then we should be able to conclude that we ought to do that thing when we consider all of the alternatives together.

Finally, suppose I decide both that I ought to go to church rather than go to the office, and that I ought to go to the office rather than go to the bar. That is, I ought to go to church out of {go to church, go to the office} and I ought to go to the office out of {go to the office, go to the bar}. Now suppose I am wondering what I ought to do out of {go to church, go to the bar}. It seems that it should just *follow* from what I've already decided that I ought to go to church rather than go to the bar: the 'ought to ... rather than ...' relation should turn out to be transitive.

But if we focus on the fact that according to contrastivism, what we ought to do can vary as the set of alternatives varies, it may be initially puzzling how we could make inferences like this. It looks as if they will be just as equivocal as the problematic inferences contrastivism blocks, which serves as the primary motivation for the theory in the first place. An 'ought' claim that is relativized to one set of alternatives wouldn't seem to tell us anything about 'ought' claims relativized to other sets of alternatives. This is how contrastivism solved the puzzles I discussed above, after all.

Fortunately, as I will now argue, the contrastivist can easily capture these inferences. Recall Jackson's claim that you ought to A out of Q just in case A is the *best* alternative in Q . The picture here is that there is a contrast-invariant ranking of actions, and relative to a set of alternatives, you ought to perform the one that ranks the highest in this ranking. With this picture in mind, we can easily get the following inferences:

Subsets: If you ought to A out of Q , then you ought to A out of subsets of Q that contain A .

Unions: If you ought to A out of Q and you ought to A out of R , then you ought to A out of $Q \cup R$, provided the members of $Q \cup R$ are mutually exclusive.

Transitivity: If you ought to A out of $\{A, B\}$ and you ought to B out of $\{B, C\}$, then you ought to A out of $\{A, C\}$.⁷

⁷ Arguments for the rationality of intransitive preferences will be relevant for assessing this inference schema. See, for example, [15] and [16]. Unfortunately, I do not have the space to discuss these challenges here. For arguments for the intransitivity of 'better than', which may be more directly relevant, see [17].

The proofs of these principles are straightforward. If A is highest-ranked in Q , then, since the ranking is contrast-invariant, it will also be highest-ranked in subsets of Q of which it is an element; so we get **Subsets**.⁸ If A is highest-ranked in Q and highest-ranked in R , then there won't be any alternative in the union of the two which ranks higher than A ; so we get **Unions**.⁹ Finally, if you ought to A rather than B and ought to B rather than C , then A is higher ranked than B , and B is higher ranked than C . Since this ranking is contrast-invariant, it just follows that A is higher ranked than C ; so we get **Transitivity**.¹⁰

So we have seen that, though contrastivism does face the challenge of accounting for cross-context deontic reasoning, the theory also has resources to explain some of them; the theory has an important kind of *structure*.

4 Inheritance-like Inferences

The inferences I discussed in the last section have to do with one kind of variation that sets of alternatives can display, which is the kind that Jackson focuses on in arguing for contrastivism. This is what I have elsewhere called *non-exhaustivity*.¹¹ Sets of alternatives can vary in terms of which possibilities they include. A subset of Q (generally) covers fewer possibilities than Q , and the union of Q and R (generally) covers more possibilities than either Q or R . What we've seen, then, is that contrastivism can account for some attractive cross-context inferences in which the sets of alternatives vary along the dimension of non-exhaustivity.

There is a second way in which sets of alternatives can vary, however. This is what we can, following [21, 5], call *resolution-sensitivity*. The sets of alternatives to which 'ought' claims are relativized can lump together the possibilities in more or less fine-grained ways, or at higher and lower resolutions. For example, the set {go to church, stay home} and the set {drive to church, take the bus to church, stay home and clean, stay home and watch football} vary in this way—the second makes distinctions between possibilities that the former does not; it divides up the relevant possibilities at a higher resolution.

This is a very important feature of contrastivism. First, in terms of general motivations for the theory, this is the feature that Cariani focuses on in develop-

⁸ It follows as a corollary of **Subsets** that if you ought to A out of Q and you ought to A out of R , then you ought to A out of the intersection, $Q \cap R$, since this will be a subset of Q that contains A . I motivated the **Unions** inference by appealing to joint deontic reasoning; you might think these kinds of consideration would instead (or sometimes) motivate this intersections principle, instead.

⁹ The restriction on **Unions**—that the members of $Q \cup R$ must be mutually exclusive—screens off cases in which there are members of Q and R that are each inconsistent with A , but not with each other (e.g., 'buy milk' and 'buy a liquid', where A is 'buy a banana'). These cases will be more complicated, but I will not discuss them here.

¹⁰ Compare the discussion of choice functions in [16] and [18]. The **Subsets** inference is similar to "basic contraction consistency" (Property α), while the **Unions** property is similar to "basic expansion consistency" (Property γ).

¹¹ See [19, 20].

ing his contrastivist semantics. Second, in terms of the purposes of this paper, some important cross-context **Inheritance**-like reasoning depends on making inferences between ‘ought’ claims that are relativized to sets of alternatives that vary along this dimension.

4.1 Contrastivism and Inheritance

It will be helpful to see exactly why **Inheritance** is problematic on a contrastivist theory. A naive contrastivist construal of this principle is the following:

Contrastivist Inheritance: If p entails q , then if it ought to be that p out of Q , it ought to be that q out of Q .

The problem here is that this doesn’t make sense on a contrastive theory, at least not for one of the simple sort we have discussed so far. This is because the alternatives in a set of alternatives are *mutually exclusive*. But if p entails q , then obviously p and q are not mutually exclusive; hence, they cannot be members of the same set of alternatives. So if it ought to be that p out of Q , then we know that q is not even in Q , so q cannot be the *best* in the set, so it cannot be true that it ought to be that q out of Q .

4.2 Nice Results

This failure of **Inheritance** on a contrastivist theory delivers some nice results. First, we can avoid Ross’s paradox (see [22]). From ‘You ought to mail the letter’, it would follow by **Inheritance** that ‘You ought to mail the letter or burn it’, and this latter claim seems false, even when the first is true. But ‘mail the letter’ and ‘mail the letter or burn it’ cannot be members of the same set of alternatives, since they are not mutually exclusive. Thus, Ross’s inference must involve a shift in the set of alternatives, and so an equivocation.

Second, we can get attractive results in Jackson and Pargetter’s Professor Procrastinate case.¹² Professor Procrastinate is asked to review a book, since she is the most qualified person to write it. However, she is a terrible procrastinator, so if she accepts, she is unlikely to actually write it—it isn’t impossible that she write it, she just probably won’t. If she accepts and does not write it, then the book will go unreviewed, which would be the worst outcome. If she declines, someone else will review it. The review won’t be as good as it would be if she did write it, but it will definitely get done.

In this case, both of the following claims seem true:

- (6) Procrastinate ought to accept and write.
- (7) It’s not the case that Procrastinate ought to accept.

(6) seems true, since Procrastinate is the most qualified person to write the review; (7) seems true since she would probably not actually get around to writing it, were she to accept.

¹² The case is originally presented in [3], and was the focal point of [23].

The problem is that accepting and writing entails accepting, so by **Inheritance**, (6) entails the negation of (7). The contrastivist avoids this contradiction by blocking this inference. ‘Accept and write’ and ‘accept’ cannot be members of the same set of alternatives, since they are not mutually exclusive. Rather, they are members of two different sets of alternatives that differ in resolution. So to infer the negation of (7), properly relativized, from (6), properly relativized, is to equivocate.

The contrastivist solution can also explain why both (6) and (7) are true, not just why they are actually consistent: (6) is relativized to {accept and write, accept and not write, not accept} and (7) to {accept, not accept}. Out of the first set, accepting and writing is the best alternative; out of the second, accepting is not the best alternative, since it would most likely lead to the worst outcome.

Note that the solutions to these puzzles depend on the resolution-sensitivity of ‘ought’. {mail the letter, burn the letter, leave it on the table} and {mail the letter or burn the letter, leave it on the table} vary in resolution—the first makes distinctions that the second does not. Similarly for {accept and write, accept and not write, not accept} and {accept, not accept}. A fine-grained ‘ought’ claim can be true while a corresponding coarse-grained ‘ought’ claim—one in which the relevant action is a coarse-grained action that “subsumes” the original fine-grained action—is false. This is perfectly fine for the contrastivist, since the ‘ought’ claims are relativized to different sets of alternatives.

4.3 Coarsening Inferences

The failure of **Inheritance** on the contrastive theory looks like a nice feature, since it provides solutions to Ross’s paradox and the Professor Procrastinate puzzle. But as should be clear by now, this also creates problems. Some **Inheritance**-like inferences look like good deontic reasoning, so we need to recapture them.

Cariani (in [5]) calls these **Inheritance**-like inferences *coarsening* inferences, since they involve moving from a more fine-grained ‘ought’ claim to a more coarse-grained ‘ought’ claim. Ross’s paradox and the Professor Procrastinate puzzle are cases in which coarsening an alternative leads to a false ‘ought’ claim. Though ‘You ought to mail the letter’ is true, coarsening to ‘mail the letter or burn it’ leads to falsity: it is not true that you ought to mail the letter or burn it. Similarly, though ‘Procrastinate ought to accept and write’ is true, coarsening to ‘accept’ leads to falsity.

But lots of coarsening inferences look like good deontic reasoning. Consider Cariani’s example:

- (8) You ought to feed your pets sufficient amounts of non-poisonous food.
- (9) You ought to feed your pets.

This looks like a fine inference to make; often coarse-grained ‘ought’ claims like (9) are perfectly acceptable, and are in some sense supported by more fine-grained ‘ought’ claims like (8). This can be so even when there are impermissible

ways of carrying out the coarse-grained option, incompatible with the more fine-grained option you ought to perform. So even though there are impermissible ways of feeding your pets—ways that are incompatible with feeding them sufficient amounts of non-poisonous food—(9) can still be true. And **Inheritance** offers a straightforward explanation of this: (8) is true, and by **Inheritance**, it entails (9).

So here is where we stand. We want to allow some coarsening inferences, like that between (8) and (9), but not too many, as in Ross’s paradox and the Professor Procrastinate puzzle. We have seen that **Inheritance** could offer an explanation of good coarsening inferences, but not without allowing bad coarsening inferences. And we have seen that contrastivism blocks the bad coarsening inferences; the danger is that it will also block the good coarsening inferences. In the next section I will introduce a sophisticated contrastivist account recently developed by Cariani that is designed to solve this problem.

5 Cariani’s Semantics

Cariani develops a version of contrastivism that is meant to walk this fine line, blocking the bad **Inheritance**-like inferences while allowing the good ones. Besides the deliberative question, which provides the alternatives, he posits two pieces of contextually-provided machinery:

- A ranking of the alternatives in the set
- A threshold in the ranking that distinguishes permissible from impermissible options; this is the *benchmark*

We then say that the agent ought to perform some action A relative to the set of alternatives Q just in case (i) all the top-ranked options are ways of A -ing, and (ii) all the ways of A -ing are permissible, i.e. above the benchmark.

The diagnosis of Ross’s paradox and the Professor Procrastinate puzzle is that there are relevant impermissible ways of performing the coarse-grained action, but not the fine-grained action. So though there are not relevant impermissible ways of mailing the letter, there are relevant impermissible ways of mailing the letter or burning it—namely, burning it. And though there are not relevant impermissible ways of accepting and writing, there are relevant impermissible ways of accepting—namely, accepting and not writing. So the coarse-grained ‘ought’ claims come out false on Cariani’s semantics.

So the problematic coarsening inferences do not go through, on Cariani’s picture. On the other hand, when we have a coarsening inference that looks like good deontic reasoning, this is because there are *not* relevant impermissible ways of performing the coarse-grained alternative. So, for (8) and (9), we can imagine that the relevant set of alternatives is {feed your pets sufficient amounts of non-poisonous food, don’t feed your pets}. Relative to this set, (8) is true, since feeding your pets sufficient amounts of non-poisonous food is the best option. (9) is also true, since (i) all the top-ranked options are ways of feeding your pets, and (ii) all the relevant ways of feeding your pets are above the threshold.

Cariani’s contrastivist picture resembles a proposal for restricting the **Inheritance** principle that Goble considers, but rejects.¹³ This restriction is to say that, if A entails B and B is permissible, if it ought to be that A , it ought to be that B . Goble rejects this principle because, when paired with the possibility of deontic conflicts, it leads to problematic results. But Cariani is not interested in accommodating deontic conflicts, and in fact they seem to be ruled out by condition (i) of his semantics, plus the assumption that the alternatives in Q are mutually exclusive.¹⁴

Though Cariani’s semantics delivers nice results in an elegant way in these cases, it also delivers some unintuitive results. For a catalogue of such problems, see [24]. Here I will just walk through two of Dowell and Bronfman’s examples involving the Professor Procrastinate case. The first targets Cariani’s claim that all the top-ranked options must be ways of A -ing, for A -ing to be what you ought to do. The second targets the claim that no relevant ways of A -ing can fall below the permissibility threshold, for A -ing to be what you ought to do.

First, in the Professor Procrastinate case, we not only want both (6) and (7) to be true. We also think that (10) is true:

(10) Procrastinate ought to not accept.

But Cariani’s semantics cannot deliver this result, since it is not true that all the top-ranked options in are ways of not accepting.

Second, consider a variation of the Procrastinate case in which Procrastinate (despite her name) is actually very reliable, and if she accepts she is very likely to write. But, crucially, it is still possible that she won’t. And given that it is pretty important that the review is written, accepting and not writing, unlikely as it is, is still plausibly a relevant alternative. Nevertheless, given that she is very likely to write, intuitively (7) is false, in this variation—Procrastinate ought to accept. But Cariani’s semantics cannot deliver this result, since there is a relevant way of accepting, namely accepting and not writing, which falls below the permissibility threshold.

In the next section I will return to the simpler contrastivist picture I introduced in sections 2 and 3 above, and show that it can deliver the right results in all of these cases.

6 A Different Picture

Return to the simpler contrastive theory I introduced above, according to which you ought to A out of Q just in case A is the highest-ranked (or best) alternative

¹³ [6], note 49.

¹⁴ Goble himself accepts a different restriction: If A entails B and A is permissible, then if it ought to be that A , it ought to be that B . As Goble shows, this principle has several nice features in a logic that allows for deontic conflicts. But notice that it also generates the **Inheritance**-based puzzles I am concerned with in this paper. For example, accepting and writing is permissible, for Professor Procrastinate, but what we want to avoid is the conclusion that since she ought to accept and write, she ought to accept. Goble’s principle validates this inference.

in Q . If we are interested in coarsening inferences, we need to think about the relationship between fine-grained alternatives, like ‘drive to the store’, and more coarse-grained alternatives, like ‘go to the store’. One way to think about this is as the relationship between an option and a disjunction of which it is a disjunct. The disjunction corresponds to a coarse-grained alternative that lumps together several relevant fine-grained ways of carrying out that alternative. So ‘go to the store’ may be identified with ‘drive to the store or walk to the store or take the bus to the store’, if those are the relevant fine-grained ways of going to the store.

So we need to think about the relationships between the value of a disjunct and the value of the disjunction of which it is a part; on the picture I’m developing here, this is the relationship between the place in the ranking of the disjuncts and the place in the ranking of the disjunction. As I will show below, there are various options here. But there are constraints that any choice will have to meet. To see some of these, consider the following overwhelmingly plausible principle:

Disjunctions: If B is better than C , then $B \vee C$ is no better than B and no worse than C — $B \vee C$ is ranked somewhere between (inclusive) B and C .¹⁵

Fortunately, this principle is just what we need to get some attractive inferences between ‘ought’ claims that are relativized to sets that differ in resolution.

6.1 Disjunctions-supported Inferences

First, **Disjunctions** gives us the following claims about the value of alternatives:

Better Dis: If A is better than both B and C , then A is better than $B \vee C$.

Worse Dis: If A is worse than both B and C , then A is worse than $B \vee C$.

Assuming the simple contrastivist claim that you ought to A out of Q just in case A is the best (or highest ranked) alternative in Q , these give us several deontic inferences, including the following three:

Dis 1: If you ought to A out of $\{A, B, C\}$, then you ought to A out of $\{A, B \vee C\}$.

Dis 2: If you ought to B out of $\{A, B\}$ and you ought to C out of $\{A, C\}$, then you ought to $B \vee C$ out of $\{A, B \vee C\}$.

Dis 3: If you ought to B out of $\{B, E, F\}$ and you ought to C out of $\{C, E, F\}$, then you ought to $B \vee C$ out of $\{B \vee C, E \vee F\}$.

Dis 1 tells us that as long as A -ing is better than all the fine-grained ways of D -ing, you ought to A rather than D . **Dis 2** tells us that as long as all the fine-grained ways of D -ing are better than A -ing, you ought to D rather than A . And **Dis 3** tells us that as long as all the fine-grained ways of D -ing are better than all the fine-grained ways of G -ing, you ought to D rather than G . The proofs of these are straightforward. These are simple, intuitive inferences. So it is important that the contrastivist theory can capture them.

¹⁵ That is, the betterness ranking is *interpolative*; see the discussion of preference rankings in [9], p. 482.

Next I will show that these principles do let us make attractive coarsening inferences, but do not license the problematic inferences in Ross’s paradox and the Professor Procrastinate puzzle.

6.2 Application to Cases

Consider the following case, which Cariani has discussed in other work ([25]). Suppose the speed limit on this road is 50 mph. Thus, you ought to drive at or under 50 mph. As Cariani points out, we should be able to infer that you ought to drive under 100 mph; this is a coarsening inference, since driving at or under 50 is a way of driving under 100. Cariani’s semantics has trouble here, since there are impermissible ways of driving under 100—namely, driving between 50 and 100.

We want to be able to infer (12) from (11):

- (11) You ought to drive at or under 50 out of {drive at or under 50, drive between 50 and 100, drive over 100}.
- (12) You ought to drive under 100 out of {drive under 100, drive over 100}.

In this case, ‘drive under 100’ in the second set of alternatives is just the disjunction of two alternatives from the first set, ‘drive at or under 50’ and ‘drive between 50 and 100’. I assume the alternatives are ranked as follows: drive at or under 50 > drive between 50 and 100 > drive over 100. From (11), we can infer that you ought to drive at or under 50 out of {drive at or under 50, drive over 100}, just given the contrastivist idea that what you ought to do out of a set of alternatives is the best thing in that set and the principle I called **Subsets** above. And from the assumption about the ranking of the alternatives I’ve just made, we can infer that you ought to drive between 50 and 100 out of {drive between 50 and 100, drive over 100}. Now we can just apply **Dis 2** to get (12). So, though (12) did not follow *directly* from (11), given the contrastive theory I’ve developed here, it does follow given that theory and the assumption about the ordinal ranking of the alternatives.

A strength of Cariani’s semantics is that it is able to give the right results in the Ross’s paradox case and in the Professor Procrastinate case. This theory can also deliver those results. We cannot infer (14) from (13):

- (13) You ought to mail the letter out of {mail the letter, burn the letter, leave the letter on the table}.
- (14) You ought to mail the letter or burn it out of {mail the letter or burn it, leave it on the table}.

The inference fails because there are fine-grained ways of mailing the letter or burning it—namely, burning it—that are ranked below ‘leave it on the table’. So we cannot use the deontic inferences generated by **Disjunctions**.

If, for some reason, the other alternatives are even more disastrous than burning the letter, things will be different. Consider, for example, {mail the letter, burn it, poison the water supply}. In this case, it may well be true that

you ought to mail the letter or burn it out of {mail the letter or burn it, poison the water supply}. But that is as it should be. An advantage of this theory over Cariani’s is that it can block Ross’s inference without predicting that ‘You ought to mail the letter or burn it’ will always be false, simply because burning the letter is impermissible. The truth of this claim may depend on how disastrous the *other* alternatives are. This point will be even more important in the Professor Procrastinate case, as we will see shortly.

The Professor Procrastinate case is similar: there are fine-grained ways of accepting—namely, accepting and not writing—that are ranked below ‘not accept’, so the **Disjunctions**-supported inferences do not let us infer that Procrastinate ought to accept from the claim that Procrastinate ought to accept and write.

Cariani’s semantics mistakenly predicted, in the original Procrastinate case, that ‘Procrastinate ought to not accept’ is always false. This is because not all of the top-ranked alternatives in {accept and write, accept and not write, not accept} were ways of not accepting. The theory here, on the other hand, can allow that this claim is true. All that has to be true is that ‘not accept’ is ranked above ‘accept’. The theory does not predict that this will always be the case, of course. But that is as it should be. After all, in the modified Procrastinate case borrowed from Dowell and Bronfman above, in which Procrastinate is very likely to write if she accepts, we want ‘Procrastinate ought to accept’ to be true. Again, Cariani’s semantics predicts that this claim will always be false, as long as ‘accept and not write’ is a relevant alternative. According to the theory here, though, the truth of this claim, and the truth of ‘Procrastinate ought to not accept’, will depend on how ‘accept’ and ‘not accept’ rank, relative to one another. And how ‘accept’ ranks will depend on how ‘accept and write or accept and not write’ ranks, since it is just equivalent to this disjunction. We have some options for determining the ranking of this disjunction, as I will now illustrate.

6.3 Options for Other Inferences

Any proposal that is consistent with the contrastivist machinery I have developed so far will have to include **Dis 1–3**, and other **Disjunctions**-supported inferences. But this leaves a great deal of leeway in either embracing or rejecting many other inferences between sets that differ in resolution. Which inferences we accept will depend on how we determine the ranking of a disjunction from the ranking of its disjuncts.

The question of how to determine the ranking of a disjunction from the rankings of its disjuncts parallels a question that arises in Hansson’s preference-based deontic logic, of how to determine the preference ranking of a sentence from the preference ranking of “holistic alternatives” in which the sentence is true.¹⁶ Holistic alternatives, for Hansson, are complete specifications of a course of action open to an agent. We can think of these, on the contrastive picture,

¹⁶ See [9], section 6.

as maximally fine-grained actions. Hansson’s sentences correspond to relatively coarse-grained actions, on the contrastive picture. Here are three options.

First, we may let the value of $A \vee B$ be a *weighted average* of the values of A and B . The value of each disjunct is weighted by its likelihood. For example, if ‘accept and write’ is much more likely than ‘accept and not write’, as in the reliable Procrastinate variation, the value of the disjunction will be relatively high; if ‘accept and not write’ is much more likely, as in the original case, the value of the disjunction will be relatively low (holding fixed the value of each disjunct). This weighted average approach delivers the verdicts I have suggested are the intuitive ones in both the original and the modified Procrastinate cases: it is not the case that unreliable Procrastinate ought to accept, and in fact is the case that she ought *not* accept, while reliable Procrastinate ought to accept. These are “actualist” intuitions; see [23].

Second, we can adopt a *pessimistic* approach, or in more friendly terms, a *cautious* approach, and let the value of the disjunction be the value of its lowest-ranked disjunct. Then no matter how likely Procrastinate is to write, ‘accept’ is going to be ranked with ‘accept and not write’.

Finally, we can adopt a more *optimistic* approach, and let the value of the disjunction be the value of its highest-ranked disjunct. Then no matter how unlikely it is that Procrastinate writes, ‘accept’ will be ranked with ‘accept and write’. This optimistic approach delivers the “possibilist” intuition that even unreliable Procrastinate ought to accept.

In fact, this optimistic approach delivers a coherent contrastivist version of **Inheritance**. Here is a simplified version of it, where the set of alternatives has just three members:

Optimistic Inheritance: If you ought to A out of $\{A, B, C\}$, then (i) you ought to $A \vee B$ out of $\{A \vee B, C\}$ and (ii) you ought to $A \vee C$ out of $\{A \vee C, B\}$.

This principle will generate the same kinds of puzzles as **Inheritance**, of course. But, though contrastivism is motivated in large part by its ability to give a principled rejection of **Inheritance**, we see now that even if you do not share the intuitions that supported this rejection, this is no reason not to be a contrastivist. Settling debates like the one between actualists and possibilists requires settling on how to determine the ranking of a coarse-grained alternative from the ranking of the fine-grained alternatives it subsumes. As I have just shown, contrastivism is in principle compatible with various options here. This kind of neutrality on substantive ethical issues is generally taken to be a good thing in a theory of the meaning of ‘ought’.

As I said above, these three options parallel options Hansson considers in developing his preference-based deontic logic. Translating between Hansson’s logics and this contrastive picture of the meaning of ‘ought’ will not be completely straightforward, but given these similarities, it does seem to me a fruitful direction for future research.

Acknowledgements. Thanks to the referees for DEON 2014, an audience at Leeds, and Noah Friedman-Biglin for helpful comments and discussion. Thanks to Shyam Nair, Mark Schroeder, and Walter Sinnott-Armstrong for discussion of earlier work that informed this paper.

References

1. Kratzer, A.: The notional category of modality. In: Eikmeyer, H.J. and Reiser, H. (eds.), *Words, Worlds, and Contexts*, pp. 38–74. de Gruyter, Berlin (1981).
2. Sloman, A.: ‘Ought’ and ‘better’. *Mind*, 79(315):385–394 (1970).
3. Jackson, F.: On the semantics and logic of obligation. *Mind*, 94(374):177–195 (1985).
4. Snedegar, J.: Contrastive semantics for deontic modals. In: Blaauw, M. (ed.), *Contrastivism in Philosophy*, Routledge (2012).
5. Cariani, F.: ‘Ought’ and resolution semantics. *Noûs*, 47(3):534–558 (2013).
6. Goble, L.: Normative conflicts and the logic of ‘ought’. *Noûs*, 43(3):450–489 (2009).
7. Goble, L.: A logic of good, should, and would: Part I. *Journal of Philosophical Logic*, 19(2):169–199 (1990).
8. van der Torre, L.: *Reasoning about Obligations: Defeasibility in Preference-based Deontic Logic*. Dissertation, Tinbergen institute research series, Erasmus University Rotterdam (1997).
9. Hansson, S. O.: Alternative semantics for deontic logic. In: Gabbay, Horty, Parent, van der Meyden, and van der Torre (eds.), *Handbook of Deontic Logic and Normative Systems*, pp. 445–497. College Publications (2013).
10. Schaffer, J.: Closure, contrast, and answer. *Philosophical Studies*, 133(2):233–255 (2007).
11. Hamblin, C. L.: Questions. *Australasian Journal of Philosophy*, 36:159–168 (1958).
12. Groenendijk, J. and Stokhof, M.: Questions. In: van Benthem, J. and ter Meulen, A. (eds.), *Handbook of Logic*, pp. 1055–1124. Elsevier Science Publishers (1997).
13. Roberts, C.: Information structure in discourse: Towards an integrated formal theory of pragmatics. *Semantics and Pragmatics*, 5:1–16 (2012).
14. Finlay, S. and Snedegar, J.: One ought too many. *Philosophy and Phenomenological Research* (forthcoming).
15. Anand, P.: The philosophy of intransitive preference. *The Economic Journal*, 103(417):337–346 (1993).
16. Sen, A.: Internal consistency of choice. *Econometrica*, 61(3):495–521 (1993).
17. Temkin, L.: *Rethinking the Good*. Oxford: Oxford University Press (2012).
18. Sen, A.: Choice functions and revealed preference. *The Review of Economic Studies*, 38(3):307–317 (1971).
19. Snedegar, J.: Reason claims and contrastivism about reasons. *Philosophical Studies*, 166(2):231–242 (2013).
20. Snedegar, J.: Contrastive reasons and promotion. *Ethics* (forthcoming).
21. Yalcin, S.: Nonfactualism about epistemic modality. In: Egan, A. and Weatherson, B. (eds.), *Epistemic Modality*. Oxford: Oxford University Press (2011).
22. Ross, A.: Imperatives and logic. *Theoria*, 7:53–71 (1941).
23. Jackson, F. Pargetter, R.: Oughts, options, and actualism. *Philosophical Review*, 95(2):233–255 (1986).
24. Dowell, J. and Bronfman, A.: The language of reasons and oughts. In: Star, D. (ed.), *Oxford Handbook of Reasons*. Oxford: Oxford University Press (forthcoming).
25. Cariani, F.: The Semantics of ‘Ought’ and the Unity of Modal Discourse. PhD Thesis, University of California, Berkeley (2009).